arXiv:1501.02069v1 [cs.LO] 9 Jan 2015

Stateless Model Checking for TSO and PSO

Parosh Abdulla, Stavros Aronis, Mohammed Faouzi Atig,
Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas

Dept. of Information Technology, Uppsala University, Sweden

Abstract. We present a technique for efficient stateless model checking of programs that execute under the relaxed memory models TSO and PSO. The basis for our technique is a novel representation of executions under TSO and PSO, called *chronological traces*. Chronological traces induce a partial order relation on relaxed memory executions, capturing dependencies that are needed to represent the interaction via shared variables. They are optimal in the sense that they only distinguish computations that are inequivalent under the widely-used representation by Shasha and Snir. This allows an optimal dynamic partial order reduction algorithm to explore a minimal number of executions while still guaranteeing full coverage. We apply our techniques to check, under the TSO and PSO memory models, LLVM assembly produced for C/pthreads programs. Our experiments show that our technique reduces the verification effort for relaxed memory models to be almost that for the standard model of sequential consistency. In many cases, our implementation significantly outperforms other comparable tools.

1 Introduction

Verification and testing of concurrent programs is difficult, since one must consider all the different ways in which instructions of different threads can be interleaved. To make matters worse, most architectures implement *relaxed memory models*, such as TSO and PSO [31,3], which make threads interact in even more and subtler ways than by standard interleaving. For example, a processor may reorder loads and stores by the same thread if they target different addresses, or it may buffer stores in a local queue.

A successful technique for finding concurrency bugs (i.e., defects that arise only under some thread schedulings), and for verifying their absence, is *stateless model checking* (SMC) [15], also known as *systematic concurrency testing* [20,34]. Starting from a test, i.e., a way to run a program and obtain some expected result, which is terminating and threadwisely deterministic (e.g. no data-nondeterminism), SMC systematically explores the set of all thread schedulings that are possible during runs of this test. A special runtime scheduler drives the SMC exploration by making decisions on scheduling whenever such decisions may affect the interaction between threads, so that the exploration covers all possible executions and detects any unexpected test results, program crashes, or assertion violations. The technique is completely automatic, has no false positives, does not suffer from memory explosion, and can easily reproduce the concurrency bugs it detects. SMC has been successfully implemented in tools such as VeriSoft [16], CHESS [24], and Concuerror [11].

There are two main problems for using SMC in programs that run under relaxed memory models (RMM). The first problem is that already under the standard model of

sequential consistency (SC) the number of possible thread schedulings grows exponentially with the length of program execution. This problem has been addressed by *partial order reduction* (POR) techniques that achieve coverage of *all* thread schedulings, by exploring only a representative subset [33,26,14,12]. POR has been adapted to SMC in the form of *Dynamic Partial Order Reduction* (DPOR) [13], which has been further developed in recent years [28,20,18,27,32,1]. DPOR is based on augmenting each execution by a *happens-before relation*, which is a partial order that captures dependencies between operations of the threads. Two executions can be regarded as equivalent if they induce the same happens-before relation, and it is therefore sufficient to explore one execution in each equivalence class (called a *Mazurkiewicz trace* [23]). DPOR algorithms guarantee to explore at least one execution in each equivalence class, thus attaining full coverage with reduced cost. A recent optimal algorithm [1] guarantees to explore *exactly* one execution per equivalence class.

The second problem is that in order to extend SMC to handle relaxed memory models, the operational semantics of programs must be extended to represent the effects of RMM. The natural approach is to augment the program state with additional structures, e.g., store buffers in the case of TSO, that model the effects of RMM [2,4,25]. This causes blow-ups in the number of possible executions, in addition to those possible under SC. However, most of these additional executions are equivalent to some SC execution. To efficiently apply SMC to handle RMM, we must therefore extend DPOR to avoid redundant exploration of equivalent executions. The natural definition of "equivalent" under RMM can be derived from the abstract representation of executions due to Shasha and Snir [30], here called *Shasha-Snir traces*, which is often used in model checking and runtime verification [17,19,9,10,6,7]. Shasha-Snir traces consist of an ordering relation between dependent operations, which generalizes the standard happens-before relation on SC executions; indeed, under SC, the equivalence relation induced by Shasha-Snir traces coincides with Mazurkiewicz traces. It would thus be natural to base DPOR for RMM on the happens-before relation induced by Shasha-Snir traces. However, this relation is in general cyclic (due to reorderings possible under RMM) and can therefore not be used as a basis for DPOR (since it is not a partial order). To develop an efficient technique for SMC under RMM we therefore need to find a different representation of executions under RMM. The representation should define an acyclic happens-before relation. Also, the induced trace equivalence should coincide with the equivalence induced by Shasha-Snir traces.

Contribution In this paper, we show how to apply SMC to TSO and PSO in a way that achieves maximal possible reduction using DPOR, in the sense that redundant exploration of equivalent executions is avoided. A cornerstone in our contribution is a novel representation of executions under RMM, called *chronological traces*, which define a happens-before relation on the events in a carefully designed representation of program executions. Chronological traces are a succinct canonical representation of executions, in the sense that there is a one-to-one correspondence between chronological traces and Shasha-Snir traces. Furthermore, the happens-before relation induced by chronological traces is a partial order, and can therefore be used as a basis for DPOR. In particular, the Optimal-DPOR algorithm of [1] will explore exactly one execution per Shasha-Snir trace. In particular, for so-called *robust* programs that are not affected

by RMM (these include data-race-free programs), Optimal-DPOR will explore as many executions under RMM as under SC: this follows from the one-to-one correspondence between chronological traces and Mazurkiewicz traces under SC. Furthermore, robustness can itself be considered a correctness criterion, which can also be automatically checked with our method (by checking whether the number of equivalence classes is increased when going from SC to RMM).

We show the power of our technique by using it to implement an efficient stateless model checker, which for C programs with pthreads explores all executions of a test-case or a program, up to some bounded length. During exploration of an execution, our implementation generates the corresponding chronological trace. Our implementation employs the source-DPOR algorithm [1], which is simpler than Optimal-DPOR, but about equally effective. Our experimental results for analyses under SC, TSO and PSO of number of intensely racy benchmarks and programs written in C/pthreads, shows that (i) the effort for verification under TSO and PSO is not much larger than the effort for verification under SC, and (ii) our implementation compares favourably against CBMC, a state-of-the-art bounded model checking tool, showing the potential of our approach.

2 Overview of Main Concepts

This section informally motivates and explains the main concepts of the paper. To focus the presentation, we consider mainly the TSO model. TSO is relevant because it is implemented in the widely used x86 as well as SPARC architectures. We first introduce TSO and its semantics. Thereafter we introduce Shasha-Snir traces, which abstractly represent the orderings between dependent events in an execution. Since Shasha-Snir traces are cyclic, we introduce an extended representation of executions, for which a natural happens-before relation is acyclic. We then describe how this happens-before relation introduces undesirable distinctions between executions, and how our new representation of chronological traces remove these distinctions. Finally, we illustrate how a DPOR algorithm exploits the happens-before relation induced by chronological traces to explore only a minimal number of executions, while still guaranteeing full coverage.

TSO — an Introduction TSO relaxes the ordering between stores and subsequent loads to different memory locations. This can be explained operationally by equipping each thread with a *store buffer* [29], which is a FIFO queue that contains pending store operations. When a thread executes a store instruction, the store does not immediately affect memory. Instead it is delayed and enqueued in the store buffer. Nondeterministically, at some later point an *update* event occurs, dequeueing the oldest store from the store buffer and updating the memory correspondingly. Load instructions take effect immediately, without being delayed. Usually a load reads a value from memory. However, if the store buffer of the same thread contains a store to the same memory location, the value is instead taken from the store in the store buffer.

p	q
store: **x** :=1	store: **y**:=1
load: $\$r$:=**y**	load: $\$s$:=**x**

Fig. 1: A program implementing the classic idiom of Dekker's mutual exclusion algorithm.

To see why this buffering semantics may cause unexpected program behaviors, consider the small program in Fig. 1. It consists of two threads p and q. The thread p first stores 1 to the memory location **x**, and then loads the value at memory location **y** into its register $\$r$. The thread q is similar. All memory locations and registers are assumed to have initial values 0.

```
p: store: x :=1  // Enqueue store
p: load: $r:=y   // Load value 0
    q: store: y:=1   // Enqueue store
    q: update        // y = 1 in memory
    q: load: $s:=x   // Load value 0
p: update        // x = 1 in memory
```

Fig. 2: An execution of the program in Fig. 1. Notice that $\$r = \$s = 0$ at the end.

It is easy to see that under the SC semantics, it is impossible for the program to terminate in a state where both registers $\$r$ and $\$s$ hold the value 0. However, under the buffering semantics of TSO, such a final state is possible. Fig. 2 shows one such program execution. We see that the store to **x** happens at the beginning of the execution, but does not take effect with respect to memory until the very end of the execution. Thus the store to **x** and the load to **y** appear to take effect in an order opposite to how they occur in the program code. This allows the execution to terminate with $\$r = \$s = 0$.

Shasha-Snir Traces for TSO Partial order reduction is based on the idea of capturing the possible orderings between dependent operations of different threads by means of a happens-before relation. When threads interact via shared variables, two instructions are considered dependent if they access the same global variable, and at least one is a write. For relaxed memory models, Shasha and Snir [30] introduced an abstract representation of executions, here referred to as *Shasha-Snir traces*, which captures such dependencies in a natural way. Shasha-Snir traces induce equivalence classes of executions. Under sequential consistency, those classes coincide with the Mazurkiewicz traces. Under a relaxed memory model, there are also additional Shasha-Snir traces corresponding to the non-sequentially consistent executions.

A Shasha-Snir trace is a directed graph, where edges capture observed event orderings. The nodes in a Shasha-Snir trace are the executed instructions. For each thread, there are edges between each pair of subsequent instructions, creating a total order for each thread. For two instructions i and j in different threads, there is an edge $i \to j$ in a trace when i causally precedes j. This happens when j reads a value that was written by i, when i reads a memory location that is subsequently updated by j, or when i and j are subsequent writes to the same memory location. In Fig. 3 we show the Shasha-Snir trace for the execution in Fig. 2.

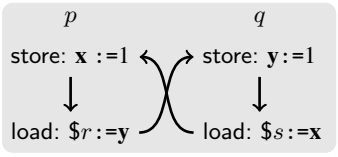

Fig. 3: The Shasha-Snir trace corresponding to the execution in Fig. 2.

Making the Happens-Before Relation Acyclic Shasha-Snir traces naturally represent the dependencies between operations in an execution, and are therefore a natural basis for applying DPOR. However, a major problem is that the happens-before relation induced by the edges is in general cyclic, and thus not a partial order. This can be seen already in the graph in Fig. 3. This problem can be addressed by adding nodes that represent explicit update events. That would be natural since such events occur in the representation of the execution in Fig. 2. When we consider the edges of the Shasha-Snir trace, we observe that although there is a conflict between p : load: $\$r$:=**y** and

q : store: $\mathbf{y}:=1$, swapping their order in the execution in Fig. 2 has no observable effect; the load still gets the same value from memory. Therefore, we should only be concerned with the order of the load relative to the update event q : update.

These observations suggest to define a representation of traces that separates stores from updates. In Fig. 4 we have redrawn the trace from Fig. 3. Updates are separated from stores, and we order updates, rather than stores, with operations of other threads. Thus, there are edges between updates to and loads from the same memory location, and between two updates to the same memory location. In Fig. 4, there is an edge from each store to the corresponding update, reflecting the principle that the update cannot occur before the store. There are edges between loads and updates of the same memory location, reflecting that swapping their order will affect the observed values. However, notice that for this program there are no edges between the updates and loads of the same thread, since they access different memory locations.

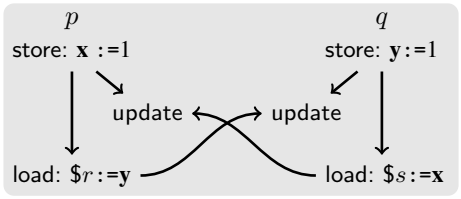

Fig. 4: A trace for the execution in Fig. 2 where updates are separated from stores.

Chronological Traces for TSO Although the new representation is a valid partial order, it will in many cases distinguish executions that are semantically equivalent according to the Shasha-Snir traces. The reason for this is the mechanism of TSO buffer forwarding: When a thread executes a load to a memory location $\mathbf{x}$, it will first check its store buffer. If the buffer contains a store to $\mathbf{x}$, then the load returns the value of the newest such store buffer entry instead of loading the value from memory. This causes difficulties for a happens-before relation that orders any update with any load of the same memory location.

For example, consider the program shown in Fig. 5. Any execution of this program will have two updates and one load to $\mathbf{x}$. Those accesses can be permuted in six different ways. Fig. 6(a), 6(b) and 6(c) show three of the corresponding happens-before relations. In Fig. 6(a) and 6(b) the load is satisfied by buffer forwarding, and in 6(c) by a read from memory. These three relations all correspond to the same Shasha-Snir trace, shown in Fig. 7(a), and they all have the same observable behavior, since the value of the load is obtained from the same store. Hence, we should find a representation of executions that does not distinguish between these three cases.

p	q
store: $\mathbf{x}:=1$	store: $\mathbf{x}:=2$
load: $\$r:=\mathbf{x}$	

Fig. 5: A program illustrating buffer forwarding.

We can now describe *chronological traces*, our representation which solves the above problems, by omitting some of the edges, leaving some nodes unrelated. More precisely, edges between loads and updates should be omitted in the following cases.

1. A load is never related to an update originating in the same thread. This captures the intuition that swapping the order of such a load and update has no effect other than changing a load from memory into a load of the same value from buffer, as seen when comparing Fig. 6(b) and 6(c).
2. A load ld from a memory location $\mathbf{x}$ by a thread p is never related to an update by an another thread q, if the update by q precedes some update to $\mathbf{x}$ originating in a

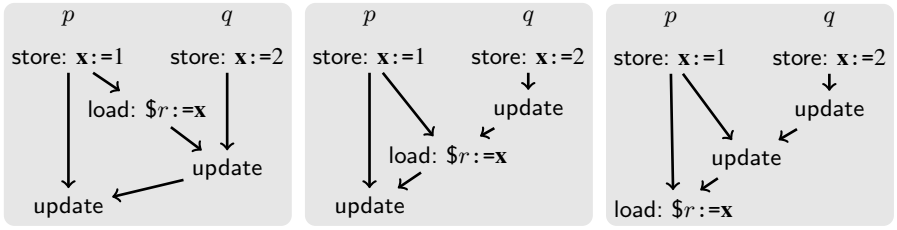

(a) (b) (c)

Fig. 6: Three redundant happens-before relations for Fig. 5.

store by p that precedes ld. This is because the value written by the update of q is effectively hidden to the load ld by the update to **x** by p. Thus, when we compare Fig. 6(a) and 6(b), we see that the order between the update by q and the load is irrelevant, since the update by q is hidden by the update by p (note that the update by p originates in a store that precedes the load).

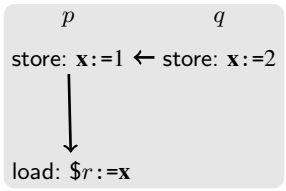

(a) A Shasha-Snir trace corresponding to all three traces of Fig. 6.

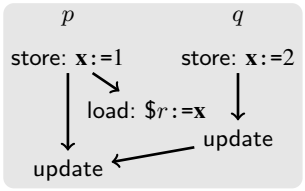

(b) The three traces can be merged into this single trace.

Fig. 7: Traces that capture all three Fig. 6(a), 6(b) and 6(c).

When we apply these rules to the example of Fig. 5, all of the three representations in Fig. 6(a), 6(b), and 6(c) merge into a single representation shown in Fig. 7(b). In total, we reduce the number of distinguished cases for the program from six to three. This is indeed the minimal number of cases that must be distinguished by any representation, since the different cases result in different values being loaded by the load instruction or different values in memory at the end of the execution. Our proposed representation is optimal for the programs in Fig. 1 and 5. In Theorem 1 of Section 3 we will show that such an optimality result holds in general.

Chronological Traces for PSO The TSO and PSO memory models are very similar. The difference is that PSO does not enforce program order between stores by the same thread to different memory locations. To capture this, chronological traces are constructed differently under TSO and PSO. In particular, under TSO there will always be edges between all updates of the same thread, but under PSO we omit those edges when

the updates access different memory locations. In Appendix C we describe in detail how to adapt the chronological traces described above to the PSO memory model.

DPOR Based on Chronological Traces Here, we illustrate how stateless model checking performs DPOR based on chronological traces, in order to explore one execution per chronological trace. As example, we use the small program of Fig. 5.

The algorithm initially explores an arbitrary execution of the program, and simultaneously generates the corresponding chronological trace. In our example, this execution can be the one shown in Fig. 8(a), along with its chronological trace. The algorithm then finds those edges of the chronological trace that can be reversed by changing the thread scheduling of the execution. In Fig. 8(a), the reversible edges are the ones from p : update to q : update, and from p : load: $\$r$:=**x** to q : update. For each such edge, the program is executed with this edge reversed. Reversing an edge can potentially lead to a completely different continuation of the execution, which must then be explored.

In the example, reversing the edge from p : load: $\$r$:=**x** to q : update will generate the execution and chronological trace in Fig. 8(b). Notice that the new execution is observably different from the previous one: the load reads the value 2 instead of 1.

The chronological traces in both Fig. 8(a) and 8(b) display a reversible edge from p : update to q : update. The algorithm therefore initiates an execution where q : update is performed before p : update. The algorithm will generate the execution and chronological trace in Fig. 8(c).

Notice that the only reversible edge in Fig. 8(c) is the one from q : update to p : update. However, executing p : update before q : update has already been explored in Fig. 8(a) and Fig. 8(b). Since there are no more edges that can be reversed, SMC terminates, having examined precisely the three chronological traces that exist for the program of Fig. 5.

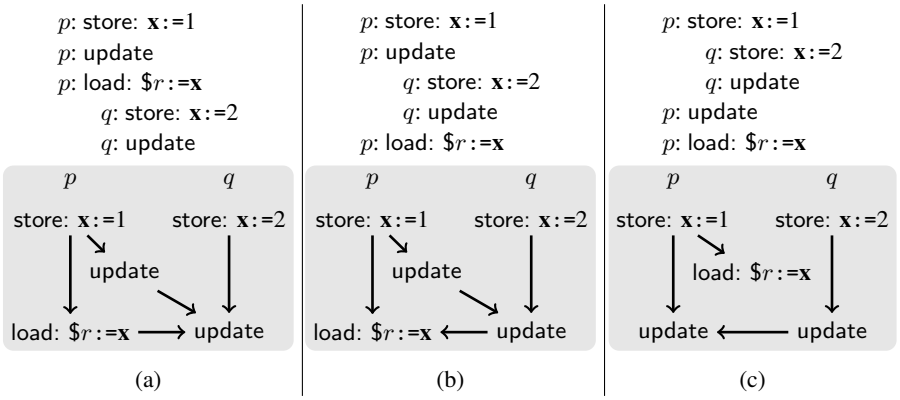

Fig. 8: How SMC with DPOR explores the program of Fig. 5.

3 Formalization

In this section we summarize our formalization of the concepts of Section 2. We introduce our representation of program executions, define chronological traces, formalize Shasha-Snir traces for TSO, and prove a one-to-one correspondence between chronological traces and Shasha-Snir traces. The formalization is self-contained, but for lack of space, we sometimes use precise English rather than formal notation. A more fully formalized version can be found in Appendix A.

Parallel Programs We consider parallel programs consisting of a number of threads that run in parallel, each executing a deterministic code, written in an assembly-like programming language. The language includes instructions store: $\mathbf{x}$:=$\$r$, load: $\$r$:=$\mathbf{x}$, and fence. Other instructions do not access memory, and their precise syntax and semantics are ignored for brevity. Here, and in the remainder of this text, **x**, **y**, **z** are used to name memory locations, u, v, w are used to name values, and $\$r$, $\$s$, $\$t$ are used to name processor registers. We use TID to denote the set of all thread identifiers.

Formal TSO Semantics We formalize the TSO model by an operational semantics. Define a *configuration* as a pair $(\mathbb{L}, \mathbb{M})$, where $\mathbb{M}$ maps memory locations to values, and $\mathbb{L}$ maps each thread p to a local configuration of the form $\mathbb{L}(p) = (\mathbb{R}, \mathbb{B})$, where $\mathbb{R}$ is the state of local registers and program counter of p, and $\mathbb{B}$ is the contents of the store buffer of p. This content is a word over pairs $(\mathbf{x}, v)$ of memory locations and values. We let the notation $\mathbb{B}(\mathbf{x})$ denote the value v such that $(\mathbf{x}, v)$ is the rightmost pair in $\mathbb{B}$ of form $(\mathbf{x}, _)$. If there is no such pair in $\mathbb{B}$, then $\mathbb{B}(\mathbf{x}) = \bot$.

In order to accommodate memory updates in our operational semantics, we assume that for each thread $p \in \mathsf{TID}$, there is an auxiliary thread $\mathsf{upd}(p)$, which nondeterministically performs memory updates from the store buffer of p. We use $\mathsf{AuxTID} = \{\mathsf{upd}(p) | p \in \mathsf{TID}\}$ to denote the set of auxiliary thread identifiers. We use p and q to refer to real or auxiliary threads in $\mathsf{TID} \cup \mathsf{AuxTID}$ as convenient.

For configurations $c = (\mathbb{L}, \mathbb{M})$ and $c' = (\mathbb{L}', \mathbb{M}')$, we write $c \xrightarrow{p} c'$ to denote that from configuration c, thread p can execute its next instruction, thereby changing the configuration into c'. Let $\mathbb{L}(p) = (\mathbb{R}, \mathbb{B})$, and $\mathbb{R}_{\mathsf{pc}}$ be obtained from $\mathbb{R}$ by advancing the program counter after p executes its next instruction. Depending on this next instruction *op*, we have the following cases.

Store: If *op* has the form store: $\mathbf{x}$:=$\$r$, then $c \xrightarrow{p} c'$ iff $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}, \mathbb{B} \cdot (\mathbf{x}, v))]$ where $v = \mathbb{R}(\$r)$, i.e., instead of updating the memory, we insert the entry $(\mathbf{x}, v)$ at the end of the store buffer of the thread.

Load: If *op* has the form load: $\$r$:=$\mathbf{x}$, then $\mathbb{M}' = \mathbb{M}$ and either

1. (**From memory**) $\mathbb{B}(\mathbf{x}) = \bot$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}[\$r \hookleftarrow \mathbb{M}(\mathbf{x})], \mathbb{B})]$, i.e., there is no entry for **x** in the thread's own store buffer, so the value is read from memory, or
2. (**Buffer forwarding**) $\mathbb{B}(\mathbf{x}) \neq \bot$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}[\$r \hookleftarrow \mathbb{B}(\mathbf{x})], \mathbb{B})]$, i.e., p reads the value of **x** from its *latest* entry in its store buffer.

Fence: If *op* has the form fence, then $c \xrightarrow{p} c'$ iff $\mathbb{B} = \varepsilon$ and $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}, \mathbb{B})]$. A fence can only be executed when the store buffer of the thread is empty.

Update: In addition to instructions which are executed by the threads, at any point when a store buffer is non-empty, an *update* event may nondeterministically occur. The memory is then updated according to the oldest (leftmost) letter in the store buffer, and that letter is removed from the buffer. To formalize this, we will assume that the auxiliary thread $\mathsf{upd}(p)$ executes a pseudo-instruction $\mathsf{u}(\mathbf{x})$. We then say that $c \xrightarrow{\mathsf{upd}(p)} c'$ iff $\mathbb{B} = (\mathbf{x}, v) \cdot \mathbb{B}'$ for some $\mathbf{x}$, v, $\mathbb{B}'$ and $\mathbb{M}' = \mathbb{M}[\mathbf{x} \hookleftarrow v]$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}, \mathbb{B}')]$.

Program Executions A program execution is a sequence $c_0 \xrightarrow{p_1} c_1 \xrightarrow{p_2} \cdots \xrightarrow{p_n} c_n$ of configurations related by transitions labelled by actual or auxiliary thread IDs. Since each transition of each program thread (including the auxiliary threads of form $\mathsf{upd}(q)$) is deterministic, a program run is uniquely determined by its sequence of thread IDs. We will therefore define an *execution* as a word of *events*. Each event represents a transition in the execution as a triple (p, i, j), where p is a regular or auxiliary thread executing an instruction i (which can possibly be an update), and the natural number j is such that the event is the jth event of p in the execution.

Chronological Traces We can now introduce the main conceptual contribution of the paper, viz. *chronological traces*. For an execution τ we define its chronological trace $\mathcal{T}_C(\tau)$ as a directed graph $\langle V, E \rangle$. The vertices V are all the events in τ (both events representing instructions and events representing updates). The edges are the union of six relations: $E = \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{su}} \cup \rightarrow_\tau^{\mathsf{uu}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{uf}}$. These edge relations are defined as follows, for two arbitrary events $e = (p, i, j), e' = (p', i', j') \in V$:

Program Order: $e \rightarrow_\tau^{\mathsf{po}} e'$ iff $p = p'$ and $j' = j + 1$, i.e., e and e' are consecutive events of the same thread.

Store to Update: $e \rightarrow_\tau^{\mathsf{su}} e'$ iff e' is the update event corresponding to the store e.

Update to Update: $e \rightarrow_\tau^{\mathsf{uu}} e'$ iff $i = \mathsf{u}(\mathbf{x})$ and $i' = \mathsf{u}(\mathbf{x})$ for some $\mathbf{x}$, and e and e' are consecutive updates to the memory location $\mathbf{x}$.

Source: $e \rightarrow_\tau^{\mathsf{src\text{-}ct}} e'$ iff e' is a load which reads the value of the update event e, which is from a different process. Notice that this definition excludes the possibility of $p = \mathsf{upd}(p')$; a load is never src-related to an update from the same thread.

Conflict: $e \rightarrow_\tau^{\mathsf{cf\text{-}ct}} e'$ iff e' is the update that overwrites the value read by e.

Update to Fence: $e \rightarrow_\tau^{\mathsf{uf}} e'$ iff $i = \mathsf{u}(\mathbf{x})$ for some $\mathbf{x}$, and $i' = \mathsf{fence}$ and $p = \mathsf{upd}(p')$ and e is the latest update by p which occurs before e' in τ. The intuition here is that the fence cannot be executed until all pending updates of the same thread have been flushed from the buffer. Hence the updates are ordered before the fence, and the chronological trace has an edge from the last of these updates to the fence event.

Shasha-Snir Traces We will now formalize Shasha-Snir traces, and prove that chronological traces are equivalent to Shasha-Snir traces, in the sense that they induce the same equivalence relation on executions. We first recall the definition of Shasha-Snir traces. We follow the formalization by Bouajjani *et al.* [7].

First, we introduce the notion of a completed execution. We say that an execution τ is *completed* when all stores have reached memory by means of a corresponding update event. In the context of Shasha-Snir traces, we will restrict ourselves to completed executions.

For a completed execution τ, we define the Shasha-Snir trace of τ as the graph $\mathcal{T}(\tau) = \langle V, E \rangle$ where V is the set of all non-update events (p, i, j) in τ (i.e., $i \neq \mathsf{u}(\mathbf{x})$ for all $\mathbf{x}$). The edges E is the union of four relations $E = \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{st}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ss}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ss}}$, where $\rightarrow_\tau^{\mathsf{po}}$ (program order) is the same as for Chronological traces, and where, letting $e = (p, i, j)$ and $e' = (p', i', j')$:

Store Order: $e \rightarrow_\tau^{\mathsf{st}} e'$ iff i and i' are two stores, whose corresponding updates are consecutive updates to the same memory location. I.e., store order defines a total order on all the stores to each memory location, based on the order in which they reach memory.

Source: $e \rightarrow_\tau^{\mathsf{src\text{-}ss}} e'$ iff e' is a load which reads its value from e, via memory or by buffer forwarding.

Conflict: $e \rightarrow_\tau^{\mathsf{cf\text{-}ss}} e'$ iff e' is the store which overwrites the value read by e.

We are now ready to state the equivalence theorem.

Theorem 1. (Equivalence of Shasha-Snir traces and chronological traces) *For a given program $\mathcal{P}$ with two completed executions τ, τ', it holds that $\mathcal{T}(\tau) = \mathcal{T}(\tau')$ iff $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$.*

The proof is found in Appendix A.

DPOR for TSO A DPOR algorithm can exploit Chronological traces to perform stateless model checking of programs that execute under TSO (and PSO), as illustrated at the end of Section 2. The explored executions follow the semantics of TSO in Section 3. For each execution, its happens-before relation, which is the transitive closure of the edge relation $E = \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{su}} \cup \rightarrow_\tau^{\mathsf{uu}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{uf}}$ of the corresponding chronological trace, is computed on the fly. Such a computation is described in more detail in Appendix B. This happens-before relation can in principle be exploited by any DPOR algorithm to explore at least one execution per equivalence class induced by Shasha-Snir traces. We state the following theorem of correctness.

Theorem 2. (Correctness of DPOR algorithms) *The algorithms* Source-DPOR *and* Optimal-DPOR *of [1], based on the happens-before relation induced by chronological traces, explore at least one execution per equivalence class induced by Shasha-Snir traces. Moreover,* Optimal-DPOR *explores exactly one execution per equivalence class.*

The proof is found in Appendix B.

4 Implementation

To show the effectiveness of our techniques we have implemented a stateless model checker for C programs. The tool, called Nidhugg, is available as open source at `https://github.com/nidhugg/nidhugg`. Major design decisions have been that Nidhugg: (i) should not be bound to a specific hardware architecture and (ii) should use an existing, mature implementation of C semantics, not implement its own. Our choice was to use the LLVM compiler infrastructure [22] and work at the level of its intermediate representation (IR). LLVM IR is low-level and allows us to analyze assembly-like

but target-independent code which is produced after employing all optimizations and transformations that the LLVM compiler performs till this stage.

Nidhugg detects assertion violations and robustness violations that occur under the selected memory model. We implement the Source-DPOR algorithm from Abdulla *et al.* [1], adapted to relaxed memory in the manner described in this paper. Before applying Source-DPOR, each spin loop is replaced by an equivalent single load and assume statement. This substantially improves the performance of Source-DPOR, since a waiting spin loop may generate a huge number of improductive loads, all returning the same wrong value; all of these loads will cause races, which will cause the number of explored traces to explode. Exploration of program executions is performed by interpretation of LLVM IR, based on the interpreter lli which is distributed with LLVM. We support concurrency through the pthreads library. This is done by hooking calls to pthread functions, and executing changes to the execution stacks (adding new threads, joining, etc.) as appropriate within the interpreter.

5 Experimental Results

We have applied our implementation to several intensely racy benchmarks, all implemented in C/pthreads. They include classical benchmarks, such as Dekker's, Lamport's (fast) and Peterson's mutual exclusion algorithms. Other programs, such as indexer.c, are designed to showcase races that are hard to identify statically. Yet others (stack_safe.c) use pthread mutexes to entirely avoid races. Lamport's algorithm and stack_safe.c originate from the TACAS Competition on Software Verification (SV-COMP). Some benchmarks originate from industrial code: apr_1.c, apr_2.c, pgsql.c and parker.c.

We show the results of our tool Nidhugg in Table 1. For comparison we also include the results of two other analysis tools, CBMC [5] and goto-instrument [4], which also target C programs under relaxed memory. The techniques of goto-instrument and CBMC are described in more detail in Section 6.

All experiments were run on a machine equipped with a 3 GHz Intel i7 processor and 6 GB RAM running 64-bit Linux. We use version 4.9 of goto-instrument and CBMC. The benchmarks have been tweaked to work for all tools, in communication with the developers of CBMC and goto-instrument. All benchmarks are available at https://github.com/nidhugg/benchmarks_tacas2015.

Table 1 shows that our technique performs well compared to the other tools for most of the examples. We will briefly highlight a few interesting results.

We see that in most cases Nidhugg pays a very modest performance price when going from sequential consistency to TSO and PSO. The explanation is that the number of executions explored by our stateless model checker is close to the number of Shasha-Snir traces, which increases very modestly when going from sequential consistency to TSO and PSO for typical benchmarks. Consider for example the benchmark stack_safe.c, which is robust, and therefore has equally many Shasha-Snir traces (and hence also chronological traces) under all three memory models. Our technique is able to benefit from this, and has almost the same run time under TSO and PSO as under SC.

			CBMC			goto-instrument			Nidhugg		
	fence	LB	SC	TSO	PSO	SC	TSO	PSO	SC	TSO	PSO
apr_1.c	-	5	t/o	t/o	t/o	t/o	!	!	**5.88**	**6.06**	**16.98**
apr_2.c	-	5	t/o	t/o	t/o	!	!	!	**2.60**	**2.20**	**5.39**
dcl_singleton.c	-	7	5.95	31.47	*18.01	5.33	5.36	*0.18	**0.08**	**0.08**	***0.08**
dcl_singleton.c	pso	7	5.88	30.98	29.45	5.20	5.18	5.17	**0.08**	**0.08**	**0.08**
dekker.c	-	10	2.42	*3.17	*2.84	1.68	*4.00	*220.11	**0.10**	***0.11**	***0.09**
dekker.c	tso	10	2.39	5.65	*3.51	1.62	297.62	t/o	**0.11**	**0.12**	***0.08**
dekker.c	pso	10	2.55	5.31	4.83	1.72	428.86	t/o	**0.11**	**0.12**	**0.12**
fib_false.c	-	-	*1.63	*3.38	*3.00	***1.60**	***1.58**	***1.56**	*2.39	*5.57	*6.20
fib_false_join.c	-	-	*0.98	*1.10	*1.91	*1.31	*0.88	*0.80	***0.32**	***0.62**	***0.71**
fib_true.c	-	-	**6.28**	9.39	7.72	6.32	**7.63**	**7.62**	25.83	75.06	86.32
fib_true_join.c	-	-	6.61	8.37	10.81	7.09	5.94	5.92	**1.20**	**2.88**	**3.19**
indexer.c	-	5	193.01	210.42	214.03	191.88	70.42	69.38	**0.10**	**0.09**	**0.09**
lamport.c	-	8	7.78	*11.63	*10.53	6.89	t/o	t/o	**0.08**	***0.08**	***0.08**
lamport.c	tso	8	7.60	26.31	*15.85	6.80	513.67	t/o	**0.09**	**0.08**	***0.07**
lamport.c	pso	8	7.72	30.92	27.51	7.43	t/o	t/o	**0.08**	**0.08**	**0.08**
parker.c	-	10	12.34	*91.99	*86.10	11.63	~~9.70~~	~~9.65~~	**1.50**	***0.09**	***0.08**
parker.c	pso	10	12.72	141.24	166.75	11.76	10.66	10.64	**1.50**	**1.92**	**2.94**
peterson.c	-	-	0.35	*0.38	*0.35	0.18	*0.20	*0.21	**0.07**	***0.07**	***0.07**
peterson.c	tso	-	0.35	0.39	*0.35	0.19	0.18	~~0.56~~	**0.07**	**0.07**	***0.07**
peterson.c	pso	-	0.35	0.41	0.40	0.18	0.18	0.19	**0.07**	**0.07**	**0.08**
pgsql.c	-	8	19.80	60.66	*4.63	21.03	46.57	*296.77	**0.08**	**0.07**	***0.08**
pgsql.c	pso	8	23.93	71.15	121.51	19.04	t/o	t/o	**0.07**	**0.07**	**0.08**
pgsql_bnd.c	pso	(4)	**3.57**	**9.55**	**12.68**	3.59	t/o	t/o	89.44	106.04	112.60
stack_safe.c	-	-	44.53	516.01	496.36	45.11	42.39	42.50	**0.34**	**0.36**	**0.43**
stack_unsafe.c	-	-	*1.40	*1.87	*2.08	*1.00	*0.81	*0.79	***0.08**	***0.08**	***0.09**
szymanski.c	-	-	0.40	*0.44	*0.43	0.23	*0.89	*1.16	**0.07**	***0.13**	***0.07**
szymanski.c	tso	-	0.40	0.50	*0.43	0.23	0.23	~~2.48~~	**0.08**	**0.08**	***0.07**
szymanski.c	pso	-	0.39	0.50	0.49	0.23	0.24	0.24	**0.08**	**0.08**	**0.08**

Table 1: Analysis times (in seconds) for our implementation Nidhugg, as well as CBMC and goto-instrument under the SC, TSO and PSO memory models. Stars (*) indicate that the analysis discovered an error in the benchmark. A t/o entry means that the tool did not terminate within 10 minutes. An ! entry means that the tool crashed. ~~Struck-out~~ entries mean that the tool gave the wrong result. In the fence column, a dash (-) means that no fences have been added to the benchmark, a memory model indicates that fences have been (manually) added to make the benchmark correct under that and stronger memory models. The LB column shows the loop unrolling depth. Superior run times are shown in bold face.

We also see that our implementation compares favourably against CBMC, a state-of-the-art bounded model checking tool, and goto-instrument. For several benchmarks, our implementation is several orders of magnitude faster.

The effect of the optimization to replace each spin loop by a load and assume statement can be seen in the pgsql.c benchmark. For comparison, we also include the benchmark pgsql_bnd.c, where the spin loop has been modified such that Nidhugg fails to automatically replace it by an assume statement.

The only other benchmark where Nidhugg is not faster is fib_true.c. The benchmark has two threads that perform the actual work, and one separate thread that checks the correctness of the computed value, causing many races, as in the case of spin loops. We show with the benchmark fib_true_join.c that in this case, the problem can be alleviated by forcing the threads to join before checking the result.

Most benchmarks in Table 1 are small program cores, ranging from 36 to 118 lines of C code, exhibiting complicated synchronization patterns. To show that our technique is also applicable to real life code, we include the benchmarks apr_1.c and apr_2.c. They each contain approximately 8000 lines of code taken from the Apache Portable Runtime library, and exercise the library primitives for thread management, locking, and memory pools. Nidhugg is able to analyze the code within a few seconds. We notice that despite the benchmarks being robust, the analysis under PSO suffers a slowdown of about three times compared to TSO. This is because the benchmarks access a large number of different memory locations. Since PSO semantics require one store buffer per memory location, this affects analysis under PSO more than under SC and TSO.

6 Related Work

To the best of our knowledge, our work is the first to apply stateless model checking techniques to the setting of relaxed memory models; see e.g. [1] for a recent survey of related work on stateless model checking and dynamic partial order reduction techniques. There have been many works dedicated to the verification and checking of programs running under RMM (e.g., [17,19,21,2,9,10,6,7,8,35]). Some of them propose *precise* analyses for checking safety properties or robustness of finite-state programs under TSO (e.g., [2,7]). Others describe monitoring and testing techniques for programs under RMM (e.g., [9,10,21]). There are also a number of efforts to design bounded model checking techniques for programs under RMM (e.g., [35,8]) which encode the verification problem in SAT.

The two closest works to ours are those presented in [5,4]. The first of them [5] develops a bounded model checking technique that can be applied to different memory models (e.g., TSO, PSO, and Power). That technique makes use of the fact that the trace of a program under RMM can be viewed as a partially ordered set. This results in a bounded model checking technique aware of the underlying memory model when constructing the SMT/SAT formula. The second line of work reduces the verification problem of a program under RMM to verification under SC of a program constructed by a code transformation [4]. This technique tries to encode the effect of the RMM semantics by augmenting the input program with buffers and queues. This work introduces also the notion of Xtop objects. Although an Xtop object is a valid acyclic rep-

resentation of Shasha-Snir traces, it will in many cases distinguish executions that are semantically equivalent according to the Shasha-Snir traces. This is never the case for chronological traces. More details on the comparison with Xtop objects can be found in Appendix D. An extensive experimental comparison with the corresponding tools [5,4] for programs running under the TSO and PSO memory models was given in Section 5.

7 Concluding Remarks

We have presented the first technique for efficient *stateless model checking* which is aware of the underlying relaxed memory model. To this end we have introduced *chronological traces* which are novel representations of executions under the TSO and PSO memory models, and induce a happens-before relation that is a partial order and can be used as a basis for DPOR. Furthermore, we have established a strict one-to-one correspondence between chronological and Shasha-Snir traces. Nidhugg, our publicly available tool, detects bugs in LLVM assembly code produced for C/pthreads programs and can be instantiated to the SC, TSO, and PSO memory models. We have applied Nidhugg to several programs, both benchmarks and of considerable size, and our experimental results show that our technique offers significantly better performance than both CBMC and goto-instrument in many cases.

We plan to extend Nidhugg to more memory models such as Power, ARM, and the C/C++ memory model. This will require to adapt the definition chronological traces to them in order to also guarantee the one-to-one correspondence with Shasha-Snir traces.

References

1. P. Abdulla, S. Aronis, B. Jonsson, and K. Sagonas. Optimal dynamic partial order reduction. In *POPL*, pages 373–384. ACM, 2014.
2. P. A. Abdulla, M. F. Atig, Y.-F. Chen, C. Leonardsson, and A. Rezine. Counter-example guided fence insertion under TSO. In *TACAS*, volume 7214 of *LNCS*, pages 204–219. Springer, 2012.
3. S. V. Adve and K. Gharachorloo. Shared memory consistency models: A tutorial. *Computer*, 29(12):66–76, 1996.
4. J. Alglave, D. Kroening, V. Nimal, and M. Tautschnig. Software verification for weak memory via program transformation. In *ESOP*, volume 7792 of *LNCS*, pages 512–532. Springer, 2013.
5. J. Alglave, D. Kroening, and M. Tautschnig. Partial orders for efficient bounded model checking of concurrent software. In *CAV*, volume 8044 of *LNCS*, pages 141–157. Springer, 2013.
6. J. Alglave and L. Maranget. Stability in weak memory models. In *CAV*, volume 6806 of *LNCS*, pages 50–66. Springer, 2011.
7. A. Bouajjani, E. Derevenetc, and R. Meyer. Checking and enforcing robustness against TSO. In *ESOP*, volume 7792 of *LNCS*, pages 533–553. Springer, 2013.
8. S. Burckhardt, R. Alur, and M. M. K. Martin. CheckFence: checking consistency of concurrent data types on relaxed memory models. In *PLDI*, pages 12–21. ACM, 2007.
9. S. Burckhardt and M. Musuvathi. Effective program verification for relaxed memory models. In *CAV*, volume 5123 of *LNCS*, pages 107–120. Springer, 2008.

10. J. Burnim, K. Sen, and C. Stergiou. Sound and complete monitoring of sequential consistency for relaxed memory models. In *TACAS*, pages 11–25. Springer, 2011. LNCS 6605.
11. M. Christakis, A. Gotovos, and K. Sagonas. Systematic testing for detecting concurrency errors in Erlang programs. In *ICST*, pages 154–163. IEEE, 2013.
12. E. M. Clarke, O. Grumberg, M. Minea, and D. Peled. State space reduction using partial order techniques. *STTT*, 2(3):279–287, 1999.
13. C. Flanagan and P. Godefroid. Dynamic partial-order reduction for model checking software. In *POPL*, pages 110–121. ACM, 2005.
14. P. Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems - An Approach to the State-Explosion Problem*, volume 1032 of *LNCS*. Springer, 1996.
15. P. Godefroid. Model checking for programming languages using VeriSoft. In *POPL*, pages 174–186. ACM Press, 1997.
16. P. Godefroid. Software model checking: The VeriSoft approach. *Formal Methods in System Design*, 26(2):77–101, 2005.
17. A. Krishnamurthy and K. A. Yelick. Analyses and optimizations for shared address space programs. *J. Parallel Distrib. Comput.*, 38(2):130–144, 1996.
18. S. Lauterburg, R. Karmani, D. Marinov, and G. Agha. Evaluating ordering heuristics for dynamic partial-order reduction techniques. In *FASE*, pages 308–322, 2010. LNCS 6013.
19. J. Lee and D. A. Padua. Hiding relaxed memory consistency with a compiler. *IEEE Trans. Computers*, 50(8):824–833, 2001.
20. Y. Lei and R. Carver. Reachability testing of concurrent programs. *IEEE Trans. Softw. Eng.*, 32(6):382–403, 2006.
21. F. Liu, N. Nedev, N. Prisadnikov, M. T. Vechev, and E. Yahav. Dynamic synthesis for relaxed memory models. In *PLDI*, pages 429–440. ACM, 2012.
22. The LLVM compiler infrastructure. `http://llvm.org`.
23. A. Mazurkiewicz. Trace theory. In *Advances in Petri Nets*, 1986.
24. M. Musuvathi, S. Qadeer, T. Ball, G. Basler, P. Nainar, and I. Neamtiu. Finding and reproducing heisenbugs in concurrent programs. In *OSDI*, pages 267–280. USENIX, 2008.
25. S. Park and D. L. Dill. An executable specification, analyzer and verifier for RMO (relaxed memory order). In *SPAA*, pages 34–41. ACM, 1995.
26. D. Peled. All from one, one for all, on model-checking using representatives. In *CAV*, volume 697 of *LNCS*, pages 409–423. Springer, 1993.
27. O. Saarikivi, K. Kähkönen, and K. Heljanko. Improving dynamic partial order reductions for concolic testing. In *ACSD*. IEEE, 2012.
28. K. Sen and G. Agha. A race-detection and flipping algorithm for automated testing of multithreaded programs. In *Haifa Verification Conference*, pages 166–182, 2007. LNCS 4383.
29. P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen. x86-TSO: a rigorous and usable programmer's model for x86 multiprocessors. *Comm. of the ACM*, 53(7):89–97, 2010.
30. D. Shasha and M. Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. on Programming Languages and Systems*, 10(2):282–312, 1988.
31. SPARC International, Inc. *The SPARC Architecture Manual Version 9*, 1994.
32. S. Tasharofi et al. TransDPOR: A novel dynamic partial-order reduction technique for testing actor programs. In *FMOODS/FORTE*, pages 219–234, 2012. LNCS 7273.
33. A. Valmari. Stubborn sets for reduced state space generation. In *Advances in Petri Nets*, volume 483 of *LNCS*, pages 491–515. Springer, 1989.
34. C. Wang, M. Said, and A. Gupta. Coverage guided systematic concurrency testing. In *ICSE*, pages 221–230. ACM, 2011.
35. Y. Yang, G. Gopalakrishnan, G. Lindstrom, and K. Slind. Nemos: A framework for axiomatic and executable specifications of memory consistency models. In *IPDPS*. IEEE, 2004.

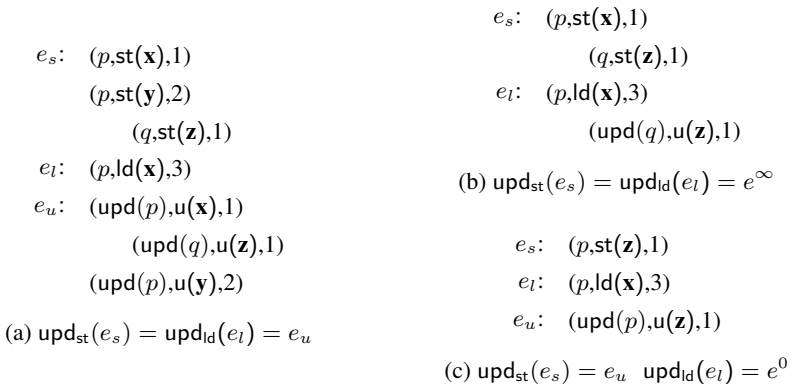

Fig. 9: Illustration of the definitions of $\mathsf{upd}_{\mathsf{st}}()$ and $\mathsf{upd}_{\mathsf{ld}}()$.

A Executions and Traces

In this appendix we introduce our representation of program executions and define chronological traces. We also formalize Shasha-Snir traces for TSO and prove that there is a one-to-one correspondence between our chronological traces and Shasha-Snir traces. For completeness, we first recall, in slightly more detail, our model of TSO and our definition of executions.

A.1 Concurrent Programs

We consider parallel programs running on shared memory under TSO. We assume that a number of threads run in parallel, each executing a deterministic code. We assume an assembly-like programming language for the code. The language includes instructions store: **x**:=$\$r$, load: $\$r$:=**x**, and fence. Other instructions do not access memory, and their precise syntax and semantics are ignored in this text for brevity. Here, and in the remainder of this text, **x**, **y**, **z** are used to name memory locations, u, v, w are used to name values, and $\$r$, $\$s$, $\$t$ are used to name registers local to processor cores. Where convenient, we will use the short forms $\mathsf{st}(\mathbf{x})$ and $\mathsf{ld}(\mathbf{x})$ to denote some store and load of **x** respectively, where the value is not interesting. We will use TID to denote the set of all thread identifiers, and MemLoc to denote the set of all memory locations.

A.2 TSO Semantics

We now define the semantics of our programming language when running under TSO. For a function f, we use the notation $f[x \hookleftarrow v]$ to denote the function f' such that $f'(x) = v$ and $f'(y) = f(y)$ whenever $y \neq x$. We use $w{\cdot}w'$ to denote the concatenation of the words w and w'.

Let the program $\mathcal{P}$ be given. We define a system configuration as a pair $(\mathbb{L}, \mathbb{M})$, where $\mathbb{L}(p)$ defines the thread local configuration for each thread p, and $\mathbb{M}(\mathbf{x})$ is the current value at $\mathbf{x}$ in memory. The local configuration of each thread is defined by $\mathbb{L}(p) = (\mathbb{R}, \mathbb{B})$. Here $\mathbb{R}$ is some structure which contains local information such as current register valuation (denoted $\mathbb{R}(\$r)$) and current program counter (denoted $\mathbb{R}(\mathsf{pc})$). $\mathbb{B}$ represents the contents of the store buffer of the thread p. It is a word over pairs $(\mathbf{x}, v)$ of memory locations and values. We let the notation $\mathbb{B}(\mathbf{x})$ denote the value v such that $(\mathbf{x}, v)$ is the rightmost letter in $\mathbb{B}$ which is of the form $(\mathbf{x}, _)$. If there is no such letter in $\mathbb{B}$, then we say that $\mathbb{B}(\mathbf{x}) = \bot$.

In order to accommodate memory updates in our operational semantics we will introduce the notion of *auxiliary threads*. For each thread $p \in \mathsf{TID}$, we assume that there is an auxiliary thread $\mathsf{upd}(p)$. The auxiliary thread $\mathsf{upd}(p)$ will nondeterministically perform memory updates from the store buffer of p, when the buffer is non-empty. We use $\mathsf{AuxTID} = \{\mathsf{upd}(p) | p \in \mathsf{TID}\}$ to denote the set of auxiliary thread identifiers. We will use p and q to refer to real or auxiliary threads in $\mathsf{TID} \cup \mathsf{AuxTID}$ as convenient.

For configurations $c = (\mathbb{L}, \mathbb{M})$ and $c' = (\mathbb{L}', \mathbb{M}')$, we use the notation $c \xrightarrow{p} c'$ to denote that from configuration c, thread p can execute its next instruction, and doing so will change the system configuration into c'. We define the transition relation $c \xrightarrow{p} c'$ depending on what the next instruction op of p is in c. In the following we assume $c = (\mathbb{L}, \mathbb{M})$ and $c' = (\mathbb{L}', \mathbb{M}')$ and $\mathbb{L}(p) = (\mathbb{R}, \mathbb{B})$. We let $\mathbb{R}_{\mathsf{pc}} = \mathbb{R}[\mathsf{pc} \hookleftarrow pc']$ where pc' is the next program counter of p after executing op.

Store: If op has the form $\mathsf{store}\colon \mathbf{x} \mathsf{:=} \$r$, then $c \xrightarrow{p} c'$ iff $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}, \mathbb{B} \cdot (\mathbf{x}, v))]$ where $v = \mathbb{R}(\$r)$. Intuitively, under TSO, instead of updating the memory with the new value v, we insert the entry $(\mathbf{x}, v)$ at the end of the store buffer of the thread.

Load: If op has the form $\mathsf{load}\colon \$r \mathsf{:=} \mathbf{x}$, then $c \xrightarrow{p} c'$ iff either

1. (**From memory**) $\mathbb{B}(\mathbf{x}) = \bot$ and $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}[\$r \hookleftarrow \mathbb{M}(\mathbf{x})], \mathbb{B})]$. Intuitively, there is no entry for $\mathbf{x}$ in the thread's own store buffer, so the value is read from memory.
2. (**Buffer forwarding**) $\mathbb{B}(\mathbf{x}) \neq \bot$ and $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}[\$r \hookleftarrow \mathbb{B}(\mathbf{x})], \mathbb{B})]$. Intuitively, we read the value of $\mathbf{x}$ from its *latest* entry in the store buffer of the thread.

Fence: If op has the form fence, then $c \xrightarrow{p} c'$ iff $\mathbb{B} = \varepsilon$ and $\mathbb{M}' = \mathbb{M}$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}_{\mathsf{pc}}, \mathbb{B})]$. A fence can be only executed when the store buffer of the thread is empty.

Update: In addition to instructions which are executed by the threads, at any point when a store buffer is non-empty, an *update* event may nondeterministically occur. The memory is then updated according to the oldest (leftmost) letter in the store buffer, and that letter is removed from the buffer. To formalize this, we will assume that the auxiliary thread $\mathsf{upd}(p)$ executes a pseudo-instruction $\mathsf{u}(\mathbf{x})$. We then say that $c \xrightarrow{\mathsf{upd}(p)} c'$ iff $\mathbb{B} = (\mathbf{x}, v) \cdot \mathbb{B}'$ for some $\mathbf{x}$, v, $\mathbb{B}'$ and $\mathbb{M}' = \mathbb{M}[\mathbf{x} \hookleftarrow v]$ and $\mathbb{L}' = \mathbb{L}[p \hookleftarrow (\mathbb{R}, \mathbb{B}')]$.

A.3 Program Executions

Based on the operational semantics in Section A.2, a program execution can be defined as a sequence $c_0 \xrightarrow{p_0} c_1 \xrightarrow{p_1} \cdots \xrightarrow{p_{n-1}} c_n$ of configurations related by transitions labelled by actual or auxiliary thread IDs. Since each transition of each program thread (including the auxiliary threads of form $\mathsf{upd}(q)$) is deterministic, a program run is uniquely determined by its sequence of thread IDs. Formally, we will therefore define each *execution* as a word of *events*. Each event is a triple (p, i, j) which represents one transition in the run. Here the thread $p \in \mathsf{TID} \cup \mathsf{AuxTID}$ can be either a regular thread $p \in \mathsf{TID}$ executing an instruction i, or an auxiliary thread $p \in \mathsf{AuxTID}$ performing an update to memory from a store buffer. In the latter case, $i = \mathsf{u}(\mathbf{x})$ is the instruction that denotes an update to a memory location. The natural number j is used to disambiguate events. We let j equal one plus the number of preceding triples (p', i', j') in the execution with $p' = p$. For an event $e = (p, i, j)$, we define $\mathsf{tid}(e) = p$. We will use Event to denote the set of all possible events. Figure 9 shows three sample executions.

For an execution τ and two events e, e' in τ, we say that $e <_\tau e'$ iff e strictly precedes e' in τ. We define two dummy events $e^0 = (\bot, \bot, -1)$ and $e^\infty = (\bot, \bot, \infty)$, and we extend $<_\tau$ such that for every event $e \notin \{e^0, e^\infty\}$ we have $e^0 <_\tau e <_\tau e^\infty$.

For an execution τ and an event $e = (p, \mathsf{st}(\mathbf{x}), j)$ in τ, we define $\mathsf{upd}_{\mathsf{st}}(e)$ to be the update event in τ corresponding to the store event e. Formally, let k be the number of events $e_w = (p', \mathsf{st}(\mathbf{y}), j')$ for any memory location $\mathbf{y}$ in τ such that $p' = p$ and $j' \leq j$. Then $\mathsf{upd}_{\mathsf{st}}(e) = (\mathsf{upd}(p), \mathsf{u}(\mathbf{x}), k)$ if there is such an event in τ. Otherwise $\mathsf{upd}_{\mathsf{st}}(e) = e^\infty$, denoting that the update is still pending at the end of τ. Figure 9(a) illustrates the typical case, where the store e_s is eventually followed by its corresponding update $\mathsf{upd}_{\mathsf{st}}(e_s) = e_u$. Figure 9(b) shows the case when the update corresponding to the store e_s is still pending at the end of the execution, and therefore $\mathsf{upd}_{\mathsf{st}}(e_s) = e^\infty$.

For an execution τ and an event $e = (p, \mathsf{ld}(\mathbf{x}), j)$ in τ, we define $\mathsf{upd}_{\mathsf{ld}}(e)$ to be the update event of the latest store to $\mathbf{x}$, which precedes e in the same thread. The intuition is that $\mathsf{upd}_{\mathsf{ld}}(e)$ is the update from which e would get its value in the case of buffer forwarding. Formally, if there is an event $e_w = (p, \mathsf{st}(\mathbf{x}), k)$ in τ such that $k < j$ and there is no event $(p, \mathsf{st}(\mathbf{x}), l)$ in τ with $k < l < j$, then $\mathsf{upd}_{\mathsf{ld}}(e) = \mathsf{upd}_{\mathsf{st}}(e_w)$. Otherwise $\mathsf{upd}_{\mathsf{ld}}(e) = e^0$. Figures 9(a) and 9(b) show the typical case, where $\mathsf{upd}_{\mathsf{ld}}(e_l)$ is taken to be the update corresponding to the latest preceding store by the same thread to the same memory location. Figure 9(c) shows the case when there is no such preceding store, and $\mathsf{upd}_{\mathsf{ld}}(e_l)$ is taken to be the dummy event e^0. (Notice that the store e_s is to a different memory location.)

A.4 Chronological Traces

We can now introduce the main conceptual contribution of the paper, viz. *chronological traces*. For an execution τ we define its chronological trace $\mathcal{T}_C(\tau)$ as a directed graph $\langle V, E \rangle$. The vertices V are all the events in τ; both events representing instructions and events representing updates. The edges are the union of six relations: $E = \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{su}} \cup \rightarrow_\tau^{\mathsf{uu}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{uf}}$.

We will illustrate the definition on an execution of the program in Figure 10(a), which contains an idiom that occurs in the mutual exclusion algorithm of Peterson. It

is mostly the same as that from Dekker's mutual exclusion algorithm. But it has two additional accesses in each thread to a separate memory location **z**. These provide an opportunity to display buffer forwarding. Figure 10(c) shows an example of an execution and Figure 10(b) shows its corresponding chronological trace.

We define the edge relations of chronological traces as follows, for two arbitrary events $e = (p, i, j), e' = (p', i', j') \in V$:

Program Order: $e \rightarrow_\tau^{\mathsf{po}} e'$ iff $p = p'$ and $j' = j + 1$. For example, in Figure 10(b) there is a program order edge from the store instruction $(p, \mathsf{st}(x), 1)$ to the store instruction $(p, \mathsf{st}(z), 2)$ which immediately follows it in the program of thread p. Similarly, the updates of each thread are program ordered. E.g., $(\mathsf{upd}(p), \mathsf{u}(\mathbf{x}), 1) \rightarrow_\tau^{\mathsf{po}} (\mathsf{upd}(p), \mathsf{u}(\mathbf{z}), 2)$.

Store to Update: $e \rightarrow_\tau^{\mathsf{su}} e'$ iff $i = \mathsf{st}(\mathbf{x})$ for some **x** and $\mathsf{upd}_{\mathsf{st}}(e) = e'$. I.e., e' is the update corresponding to the store e. This is illustrated in Figure 10(b) where there is an su-edge from each store, to its corresponding update.

Update to Update: $e \rightarrow_\tau^{\mathsf{uu}} e'$ iff $i = \mathsf{u}(\mathbf{x})$ and $i' = \mathsf{u}(\mathbf{x})$ for some **x** and $e <_\tau e'$ and there is no event $e'' = (p'', \mathsf{u}(\mathbf{x}), j'')$ such that $e <_\tau e'' <_\tau e'$. I.e., $\rightarrow_\tau^{\mathsf{uu}}$ defines the total, chronological order on updates to each memory location. In Figure 10(b) we see that the two updates to **z** are uu-ordered with each other in the same order as they appear in the execution in Figure 10(c). However, they are not uu-ordered with the updates to **x** and **y**.

Source: $e \rightarrow_\tau^{\mathsf{src\text{-}ct}} e'$ iff for some **x** it holds that $i = \mathsf{u}(\mathbf{x})$ and $i' = \mathsf{ld}(\mathbf{x})$ and $\mathsf{upd}_{\mathsf{ld}}(e') <_\tau e <_\tau e'$ and there is no update $e'' = (p'', \mathsf{u}(\mathbf{x}), j'')$ to **x** such that $e <_\tau e'' <_\tau e'$. I.e., if the source of the value read by e' is an update e from a different process, then $e \rightarrow_\tau^{\mathsf{src\text{-}ct}} e'$. Otherwise, there is no incoming $\rightarrow_\tau^{\mathsf{src\text{-}ct}}$ edge to e'. Notice that this definition excludes the possibility of $p = \mathsf{upd}(p')$; a load is never src-related to an update from the same thread. In Figure 10(b) we see that the load $(q, \mathsf{ld}(\mathbf{x}), 4)$ takes its value from the update $(\mathsf{upd}(p), \mathsf{u}(\mathbf{x}), 1)$. Therefore the events are src-related. But the loads to **z** both read the value written by their own thread, and therefore have no src-relation.

Conflict: $e \rightarrow_\tau^{\mathsf{cf\text{-}ct}} e'$ iff $i = \mathsf{ld}(\mathbf{x})$ and $i' = \mathsf{u}(\mathbf{x})$ for some **x** and e' is the first (w.r.t. $<_\tau$) event e_u of the form $(_, \mathsf{u}(\mathbf{x}), _)$ such that both $e <_\tau e_u$ and $\mathsf{upd}_{\mathsf{ld}}(e) <_\tau e_u$. The intuition here is that $e \rightarrow_\tau^{\mathsf{cf\text{-}ct}} e'$ when e' is the first update which succeeds e in the coherence order of **x**. Equivalently, e' is the update that overwrites the value that was read by e. In Figure 10(b), the load to **y** by p reads the initial value of **y**, which is then overwritten by the update to **y** by q. Therefore the load has a cf-edge to the update. The load to **z** by p reads the value of $(\mathsf{upd}(p), \mathsf{u}(\mathbf{z}), 2)$ by buffer forwarding. That value is later overwritten in memory by the update $(\mathsf{upd}(q), \mathsf{u}(\mathbf{z}), 2)$. Therefore the load has a cf-edge to the update originating in thread q.

Update to Fence: $e \rightarrow_\tau^{\mathsf{uf}} e'$ iff $i = \mathsf{u}(\mathbf{x})$ for some **x**, and $i' = \mathsf{fence}$ and $p = \mathsf{upd}(p')$ and $e <_\tau e'$ and there is no event $e'' = (p, \mathsf{u}(\mathbf{y}), j'')$ for any **y** such that $e <_\tau e'' <_\tau e'$. The intuition here is that the fence cannot be executed until all pending updates of the same thread have been flushed from the buffer. Hence the updates are ordered before the fence.

A.5 Shasha-Snir Traces

We will now prove that chronological traces are equivalent to Shasha-Snir traces, in the sense that there is a one-to-one mapping between Shasha-Snir traces T and chronological traces T_C such that the set of executions corresponding to T is the same as the set of executions corresponding to T_C.

We briefly recall the definition of Shasha-Snir traces, based on the definition by Bouajjani *et al.* [7].

First, we introduce the notion of a *completed* execution. We say that an execution τ is completed when all stores have reached memory, i.e., when for every event $e = (p, \mathsf{st}(\mathbf{x}), j)$ in τ we have $\mathsf{upd}_{\mathsf{st}}(e) \neq e^{\infty}$. In the context of Shasha-Snir traces, we will restrict ourselves to completed executions.

For a completed execution τ, we define the Shasha-Snir trace of τ as the graph $\mathcal{T}(\tau) = \langle V, E \rangle$ where V is the set of all non-update events (p, i, j) in τ where $i \neq \mathsf{u}(\mathbf{x})$ for all $\mathbf{x}$. The edges E is the union of four relations $E = \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{st}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ss}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ss}}$.

For two arbitrary events $e = (p, i, j), e' = (p', i', j') \in V$, we define the relations as follows:

Program Order: $e \rightarrow_\tau^{\mathsf{po}} e'$ iff $p = p'$ and $j' = j + 1$. This is the same program order as for chronological traces.

Store Order: $e \rightarrow_\tau^{\mathsf{st}} e'$ iff $i = \mathsf{st}(\mathbf{x})$ and $i' = \mathsf{st}(\mathbf{x})$ and the corresponding updates are ordered in τ s.t. $\mathsf{upd}_{\mathsf{st}}(e) <_\tau \mathsf{upd}_{\mathsf{st}}(e')$ and there is no other update event $e'' = (p'', \mathsf{u}(\mathbf{x}), j'')$ such that $\mathsf{upd}_{\mathsf{st}}(e) <_\tau e'' <_\tau \mathsf{upd}_{\mathsf{st}}(e')$. I.e., store order defines a total order on all the stores to each memory location, based on the order in which they reach memory.

Source: $e \rightarrow_\tau^{\mathsf{src\text{-}ss}} e'$ iff $i' = \mathsf{ld}(\mathbf{x})$ and e is the maximal store event $e'' = (p'', \mathsf{st}(\mathbf{x}), j'')$ with respect to $\rightarrow_\tau^{\mathsf{st}*}$ such that either $\mathsf{upd}_{\mathsf{st}}(e'') <_\tau e'$ or $e'' \rightarrow_\tau^{\mathsf{po}*} e'$. I.e., $e \rightarrow_\tau^{\mathsf{src\text{-}ss}} e'$ when e' is a load which reads its value from e, via memory or by buffer forwarding.

Conflict: $e \rightarrow_\tau^{\mathsf{cf\text{-}ss}} e'$ iff $i = \mathsf{ld}(\mathbf{x})$ and $i' = \mathsf{st}(\mathbf{x})$ and if there is an event e'' such that $e'' \rightarrow_\tau^{\mathsf{src\text{-}ss}} e$ then $e'' \rightarrow_\tau^{\mathsf{st}} e'$, otherwise e' has no predecessor in $\rightarrow_\tau^{\mathsf{st}}$. I.e., e' is the store which overwrites the value that was read by e.

The definition of Shasha-Snir traces is illustrated in Figure 10(d). We are now ready to state the equivalence theorem.

Theorem 1. (Equivalence of Shasha-Snir traces and chronological traces) *For a given program $\mathcal{P}$ with two completed executions τ, τ', it holds that $\mathcal{T}(\tau) = \mathcal{T}(\tau')$ iff $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$.*

Proof We decompose the theorem into the following two lemmas, which are proven separately.

Lemma 1. (Equivalence of Shasha-Snir traces and chronological traces: $\Rightarrow$ direction) *For a given program $\mathcal{P}$ with two completed executions τ, τ', it holds that if $\mathcal{T}(\tau) = \mathcal{T}(\tau')$ then $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$.*

Lemma 2. (Equivalence of Shasha-Snir traces and chronological traces: $\Leftarrow$ direction) *For a given program $\mathcal{P}$ with two completed executions τ, τ', it holds that if $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$ then $\mathcal{T}(\tau) = \mathcal{T}(\tau')$.*

Proof of Lemma 1 Let two completed executions τ and τ' be given. Let
$\mathcal{T}(\tau) = \langle V_{SS}, \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{st}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ss}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ss}} \rangle$ and
$\mathcal{T}(\tau') = \langle V'_{SS}, \rightarrow_{\tau'}^{\mathsf{po}} \cup \rightarrow_{\tau'}^{\mathsf{st}} \cup \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} \cup \rightarrow_{\tau'}^{\mathsf{cf\text{-}ss}} \rangle$ and
$\mathcal{T}_C(\tau) = \langle V_C, \rightarrow_\tau^{\mathsf{po}} \cup \rightarrow_\tau^{\mathsf{su}} \cup \rightarrow_\tau^{\mathsf{uu}} \cup \rightarrow_\tau^{\mathsf{src\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{cf\text{-}ct}} \cup \rightarrow_\tau^{\mathsf{uf}} \rangle$ and
$\mathcal{T}_C(\tau') = \langle V'_C, \rightarrow_{\tau'}^{\mathsf{po}} \cup \rightarrow_{\tau'}^{\mathsf{su}} \cup \rightarrow_{\tau'}^{\mathsf{uu}} \cup \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} \cup \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}} \cup \rightarrow_{\tau'}^{\mathsf{uf}} \rangle$.
Furthermore, assume that $\mathcal{T}(\tau) = \mathcal{T}(\tau')$.

First, we determine that the events are the same in both chronological traces: $V_C = V'_C$. From $V_{SS} = V'_{SS}$ we have that the non-update events in τ are the same as the ones in τ'. Since τ and τ' contain the same stores for each thread in the same per-thread order, it follows from the completedness of τ and τ', and from the TSO semantics that τ and τ' also have the same update events. Hence $V_C = V'_C$.

We see that the definitions of program order and store to update order in chronological traces are entirely determined by which events exist in the execution for each thread. Since both executions have the same events, we conclude that $\rightarrow_\tau^{\mathsf{po}} = \rightarrow_{\tau'}^{\mathsf{po}}$ and $\rightarrow_\tau^{\mathsf{su}} = \rightarrow_{\tau'}^{\mathsf{su}}$. The equality of update to fence order follows similarly.

Let us consider the definitions of update to update order for chronological traces and store order for Shasha-Snir traces. We see that there is a one-to-one mapping between relations $e \rightarrow_\tau^{\mathsf{st}} e'$ for stores in Shasha-Snir traces to relations $\mathsf{upd_{st}}(e) \rightarrow_\tau^{\mathsf{uu}} \mathsf{upd_{st}}(e')$ in chronological traces. Since the store orders are the same for τ and τ', we thus conclude that the update to update orders are also the same: $\rightarrow_\tau^{\mathsf{uu}} = \rightarrow_{\tau'}^{\mathsf{uu}}$.

We now turn our attention to proving that $\rightarrow_\tau^{\mathsf{src\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$. We will prove that $\rightarrow_\tau^{\mathsf{src\text{-}ct}} \subseteq \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$. From symmetry it then follows that $\rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} \subseteq \rightarrow_\tau^{\mathsf{src\text{-}ct}}$, and hence $\rightarrow_\tau^{\mathsf{src\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$. Let us assume that the relation $e \rightarrow_\tau^{\mathsf{src\text{-}ct}} e'$ exists in $\rightarrow_\tau^{\mathsf{src\text{-}ct}}$for some events $e = (p, \mathsf{u}(\mathbf{x}), j)$ and $e = (p', \mathsf{ld}(\mathbf{x}), j')$. We will prove that the same relation $e \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} e'$ exists in $\rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$. From the definition of $\rightarrow_\tau^{\mathsf{src\text{-}ct}}$ we have that $\mathsf{upd_{ld}}(e') <_\tau e <_\tau e'$ and there is no update $e'' = (p'', \mathsf{u}(\mathbf{x}), i'')$ to the same memory location such that $e <_\tau e'' <_\tau e'$. Since e' is preceded in τ by at least one update to $\mathbf{x}$, there must be a store event e_w such that $e_w \rightarrow_\tau^{\mathsf{src\text{-}ss}} e'$ in τ. From the definition of $\rightarrow_\tau^{\mathsf{src\text{-}ss}}$ we have that e_w is the maximal event $(p'', \mathsf{st}(\mathbf{x}), j'')$ with respect to $\rightarrow_\tau^{\mathsf{st}*}$ such that either $\mathsf{upd_{st}}(e_w) <_\tau e'$ or $e_w \rightarrow_\tau^{\mathsf{po}*} e'$. If $e_w \rightarrow_\tau^{\mathsf{po}*} e'$, then $\mathsf{upd_{st}}(e_w) = \mathsf{upd_{ld}}(e')$. But then the maximality of e_w contradicts $\mathsf{upd_{ld}}(e') <_\tau e <_\tau e'$. Hence we have $\mathsf{upd_{st}}(e_w) <_\tau e'$. Maximality of e_w now gives that $\mathsf{upd_{st}}(e_w) = e$. Since $\rightarrow_\tau^{\mathsf{src\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$ we have that in τ' also $e_w \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} e'$. From the definition of $\rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$ and $\neg(e_w \rightarrow_{\tau'}^{\mathsf{po}*} e')$ we know that $\mathsf{upd_{st}}(e_w)$ is the store-order-maximal update to $\mathbf{x}$ that precedes e' in τ'. Since the store order is the same for τ and τ' we have $\mathsf{upd_{ld}}(e') <_{\tau'} e$. But then $e = \mathsf{upd_{st}}(e_w)$ satisfies the criteria for $e \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} e'$.

Finally, we will show that $\rightarrow_\tau^{\mathsf{cf\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}}$. Similarly to the proof for $\rightarrow_\tau^{\mathsf{src\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$, it suffices here to show that $\rightarrow_\tau^{\mathsf{cf\text{-}ct}} \subseteq \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}}$. Assume therefore that $e_r \rightarrow_\tau^{\mathsf{cf\text{-}ct}} e_u$ for some events $e_r = (p, \mathsf{ld}(\mathbf{x}), j)$, $e_u = (p', \mathsf{u}(\mathbf{x}), j')$. We will show that $e_r \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}} e_u$. The definition of $\rightarrow_\tau^{\mathsf{cf\text{-}ct}}$ gives that e_u is the first (w.r.t. $<_\tau$) event e of the form $(_, \mathsf{u}(\mathbf{x}), _)$ such that both $e_r <_\tau e$ and $\mathsf{upd_{ld}}(e_r) <_\tau e$. Let e_w be the store event such that

$\mathsf{upd}_{\mathsf{st}}(e_w) = e_u$. We will split the proof in cases depending on whether or not there exists a source event for e_r in the Shasha-Snir traces.

Assume therefore first (i) that there is no event e_{src} such that $e_{src} \rightarrow_{\tau}^{\mathsf{src\text{-}ss}} e_r$. Then there is no update to $\mathbf{x}$ that precedes e_r in $<_\tau$. Furthermore $\mathsf{upd}_{\mathsf{ld}}(e_r) = e^0$. This tells us that e_w has no predecessor in $\rightarrow_{\tau}^{\mathsf{st}}$. Since $\rightarrow_{\tau}^{\mathsf{st}} = \rightarrow_{\tau'}^{\mathsf{st}}$, we also have that e_w has no predecessor in $\rightarrow_{\tau'}^{\mathsf{st}}$. Furthermore, since e_r has no source event in τ', it must be the case that $e_r <_{\tau'} e_u$. But then, e_u is the first update event in τ' which is after both e_r and $\mathsf{upd}_{\mathsf{ld}}(e_r)$. And so we have $e_r \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}} e_u$.

Next assume (ii) that there is an event e_{src} with $e_{src} \rightarrow_{\tau}^{\mathsf{src\text{-}ss}} e_r$ and that $\mathsf{tid}(e_{src}) = \mathsf{tid}(e_r)$. Then it must be the case that $\mathsf{upd}_{\mathsf{st}}(e_{src}) = \mathsf{upd}_{\mathsf{ld}}(e_r)$. Since $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$, we have that $e_{src} \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} e_r$. There can be no update event e to the same memory location $\mathbf{x}$ such that $\mathsf{upd}_{\mathsf{ld}}(e_r) <_\tau e <_\tau e_r$. If there were such an e, then e_{src} wouldn't be the source of e_r. The same argument goes in τ'. This tells us that e_u is the immediate store order successor of $\mathsf{upd}_{\mathsf{ld}}(e_r)$, i.e., $\mathsf{upd}_{\mathsf{ld}}(e_r) \rightarrow_{\tau}^{\mathsf{uu}} e_u$ and $e_{src} \rightarrow_{\tau}^{\mathsf{st}} e_w$. Since $\rightarrow_{\tau}^{\mathsf{uu}} = \rightarrow_{\tau'}^{\mathsf{uu}}$, we have $\mathsf{upd}_{\mathsf{ld}}(e_r) \rightarrow_{\tau'}^{\mathsf{uu}} e_u$. Hence e_u is the first update event which succeeds both e_r and $\mathsf{upd}_{\mathsf{ld}}(e_r)$ in $<_{\tau'}$. Thus $e_r \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}} e_u$.

Lastly, we assume (iii) that there is an event e_{src} such that $e_{src} \rightarrow_{\tau}^{\mathsf{src\text{-}ss}} e_r$ and that $\mathsf{tid}(e_{src}) \neq \mathsf{tid}(e_r)$. Then it is the case in τ that $\mathsf{upd}_{\mathsf{ld}}(e_r) <_\tau \mathsf{upd}_{\mathsf{st}}(e_{src}) <_\tau e_r$. And there is no update event e to $\mathbf{x}$ such that $\mathsf{upd}_{\mathsf{st}}(e_{src}) <_\tau e <_\tau e_r$. The same holds in τ'. Since e_u is the first update to $\mathbf{x}$ after e_r in τ, this means that we have $\mathsf{upd}_{\mathsf{st}}(e_{src}) \rightarrow_{\tau}^{\mathsf{uu}} e_u$. We have $\rightarrow_{\tau}^{\mathsf{uu}} = \rightarrow_{\tau'}^{\mathsf{uu}}$, so $\mathsf{upd}_{\mathsf{st}}(e_{src}) \rightarrow_{\tau'}^{\mathsf{uu}} e_u$. Now it must be the case that $e_r <_{\tau'} e_u$. Otherwise, e_{src} wouldn't be the source of e_r in τ', and we know $e_{src} \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} e_r$. Hence e_u is an update event that succeeds both e_r and $\mathsf{upd}_{\mathsf{ld}}(e_r)$ in $<_{\tau'}$. It remains to show that it is the first such update. Suppose $e \neq e_u$ is an update event to $\mathbf{x}$ such that $e_r <_\tau e <_\tau e_u$. Then it would be the case that $\mathsf{upd}_{\mathsf{st}}(e_{src}) <_{\tau'} e <_{\tau'} e_u$. But this would contradict $\mathsf{upd}_{\mathsf{st}}(e_{src}) \rightarrow_{\tau}^{\mathsf{uu}} e_u$. Thus we have $e_r \rightarrow_{\tau'}^{\mathsf{cf\text{-}ss}} e_u$.

This concludes the proof of $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$.

Proof of Lemma 2 Let two completed executions τ and τ' be given. Let
$\mathcal{T}(\tau) = \langle V_{SS}, \rightarrow_{\tau}^{\mathsf{po}} \cup \rightarrow_{\tau}^{\mathsf{st}} \cup \rightarrow_{\tau}^{\mathsf{src\text{-}ss}} \cup \rightarrow_{\tau}^{\mathsf{cf\text{-}ss}} \rangle$ and
$\mathcal{T}(\tau') = \langle V'_{SS}, \rightarrow_{\tau'}^{\mathsf{po}} \cup \rightarrow_{\tau'}^{\mathsf{st}} \cup \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} \cup \rightarrow_{\tau'}^{\mathsf{cf\text{-}ss}} \rangle$ and
$\mathcal{T}_C(\tau) = \langle V_C, \rightarrow_{\tau}^{\mathsf{po}} \cup \rightarrow_{\tau}^{\mathsf{su}} \cup \rightarrow_{\tau}^{\mathsf{uu}} \cup \rightarrow_{\tau}^{\mathsf{src\text{-}ct}} \cup \rightarrow_{\tau}^{\mathsf{cf\text{-}ct}} \cup \rightarrow_{\tau}^{\mathsf{uf}} \rangle$ and
$\mathcal{T}_C(\tau') = \langle V'_C, \rightarrow_{\tau'}^{\mathsf{po}} \cup \rightarrow_{\tau'}^{\mathsf{su}} \cup \rightarrow_{\tau'}^{\mathsf{uu}} \cup \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} \cup \rightarrow_{\tau'}^{\mathsf{cf\text{-}ct}} \cup \rightarrow_{\tau'}^{\mathsf{uf}} \rangle$.
Furthermore, assume that $\mathcal{T}_C(\tau) = \mathcal{T}_C(\tau')$.

We will prove that $\mathcal{T}(\tau) = \mathcal{T}(\tau')$. We know that V_{SS} (respectively V'_{SS}) is precisely the non-updates of V_C (respectively V'_C). Since $V_C = V'_C$ we have $V_{SS} = V'_{SS}$.

For the relations $\rightarrow_{\tau}^{\mathsf{po}}$ and $\rightarrow_{\tau}^{\mathsf{st}}$, a reasoning analogue to that in the $\Rightarrow$ direction gives that $\rightarrow_{\tau}^{\mathsf{po}} = \rightarrow_{\tau'}^{\mathsf{po}}$ and $\rightarrow_{\tau}^{\mathsf{st}} = \rightarrow_{\tau'}^{\mathsf{st}}$.

We will show that $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}} \subseteq \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$. Symmetry then gives $\rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} \subseteq \rightarrow_{\tau}^{\mathsf{src\text{-}ss}}$, and hence $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$. Assume therefore that $e_w \rightarrow_{\tau}^{\mathsf{src\text{-}ss}} e_r$ holds for some events $e_w = (p, \mathsf{st}(\mathbf{x}), j)$ and $e_r = (p', \mathsf{ld}(\mathbf{x}), j')$. Then by the definition of $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}}$ we have that e_w is the maximal event $e = (p'', \mathsf{st}(\mathbf{x}), j'')$ with respect to $\rightarrow_{\tau}^{\mathsf{st}*}$ such that either $\mathsf{upd}_{\mathsf{st}}(e) <_\tau e_r$ or $e \rightarrow_{\tau}^{\mathsf{po}*} e_r$. We will separate the proof by cases: either $\mathsf{tid}(e_w) = \mathsf{tid}(e_r)$ or $\mathsf{tid}(e_w) \neq \mathsf{tid}(e_r)$.

Assume first (i) that $\mathsf{tid}(e_w) = \mathsf{tid}(e_r)$. Then it holds that $e_w \rightarrow_{\tau}^{\mathsf{po}*} e_r$, since the events must be program ordered, and the other direction implies $e_r <_\tau \mathsf{upd}_{\mathsf{st}}(e_w)$.

Program order is the same in τ' as in τ, so we also have $e_w \rightarrow_{\tau'}^{\mathsf{po}*} e_r$. It remains to show that e_w is maximal in τ'. First we conclude that there can be no store event e such that $e_w \rightarrow_{\tau'}^{\mathsf{st}} e$ and $e \rightarrow_{\tau'}^{\mathsf{po}*} e_r$. This is because both the program order and the store order are the same in τ' as in τ, and hence such an event e would contradict the assumed maximality of e_w w.r.t. τ. As a corollary we have $\mathsf{upd}_{\mathsf{ld}}(e_r) = \mathsf{upd}_{\mathsf{st}}(e_w)$. Next we need to conclude that there is no event e such that $e_w \rightarrow_{\tau'}^{\mathsf{st}} e$ and $\mathsf{upd}_{\mathsf{st}}(e) <_{\tau'} e_r$. We know that there is no such event in τ: i.e., there is no event e such that $e_w \rightarrow_{\tau}^{\mathsf{st}} e$ and $\mathsf{upd}_{\mathsf{st}}(e) <_{\tau} e_r$. Hence by the definition of $\rightarrow_{\tau}^{\mathsf{src\text{-}ct}}$ there is no event e_{src}^C which is source related with e_r in the chronological trace: $e_{src}^C \rightarrow_{\tau}^{\mathsf{src\text{-}ct}} e_r$. Since $\rightarrow_{\tau}^{\mathsf{src\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$, the same holds in τ'. Now if there were an event such as e in τ', then e_r would have a source according to $\rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$. This is a contradiction, and so there can be no such e in τ'. Hence, e_w is the maximal store event w.r.t. $\rightarrow_{\tau'}^{\mathsf{st}\,*}$ which is either updated $<_{\tau'}$-before e_r or program order-before e_r. That concludes the proof for the case that $\mathsf{tid}(e_w) = \mathsf{tid}(e_r)$.

Next assume (ii) that $\mathsf{tid}(e_w) \neq \mathsf{tid}(e_r)$. Clearly e_w is not program ordered with e_r. Hence the definition of $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}}$ gives that $\mathsf{upd}_{\mathsf{st}}(e_w) <_{\tau} e_r$. The maximality of e_w gives that $\mathsf{upd}_{\mathsf{ld}}(e_r) <_{\tau} \mathsf{upd}_{\mathsf{st}}(e_w)$, and that there is no update event $e = (p'', \mathsf{u}(\mathbf{x}), j'')$ such that $\mathsf{upd}_{\mathsf{st}}(e_w) <_{\tau} e <_{\tau} e_r$. Then we have $\mathsf{upd}_{\mathsf{st}}(e_w) \rightarrow_{\tau}^{\mathsf{src\text{-}ct}} e_r$ by the definition of $\rightarrow_{\tau}^{\mathsf{src\text{-}ct}}$. By $\rightarrow_{\tau}^{\mathsf{src\text{-}ct}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$ we also have $\mathsf{upd}_{\mathsf{st}}(e_w) \rightarrow_{\tau'}^{\mathsf{src\text{-}ct}} e_r$. By the definition of $\rightarrow_{\tau'}^{\mathsf{src\text{-}ct}}$ we now have that e_w is the greatest (w.r.t. $<_{\tau'}$) store event with $\mathsf{upd}_{\mathsf{st}}(e_w) <_{\tau'} e_r$. We also have that $\mathsf{upd}_{\mathsf{ld}}(e_r) <_{\tau'} \mathsf{upd}_{\mathsf{st}}(e_w)$. Since there can be no event $e = (_, \mathsf{st}(\mathbf{x}), _)$ such that $e \rightarrow_{\tau'}^{\mathsf{po}*} e_r$ and $\mathsf{upd}_{\mathsf{ld}}(e_r) <_{\tau'} \mathsf{upd}_{\mathsf{st}}(e)$, we have that e_w is the maximal event $e = (_, \mathsf{st}(\mathbf{x}), _)$ with respect to $\rightarrow_{\tau}^{\mathsf{st}\,*}$ such that either $\mathsf{upd}_{\mathsf{st}}(e) <_{\tau'} e_r$ or $e \rightarrow_{\tau'}^{\mathsf{po}*} e_r$. Hence $e_w \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}} e_r$. This concludes the proof for $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$.

Since $\rightarrow_{\tau}^{\mathsf{cf\text{-}ss}}$ (respectively $\rightarrow_{\tau'}^{\mathsf{cf\text{-}ss}}$) is entirely determined by $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}}$ and $\rightarrow_{\tau}^{\mathsf{st}}$ (respectively $\rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$ and $\rightarrow_{\tau'}^{\mathsf{st}}$), and we know that $\rightarrow_{\tau}^{\mathsf{src\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{src\text{-}ss}}$ and $\rightarrow_{\tau}^{\mathsf{st}} = \rightarrow_{\tau'}^{\mathsf{st}}$, we immediately get that $\rightarrow_{\tau}^{\mathsf{cf\text{-}ss}} = \rightarrow_{\tau'}^{\mathsf{cf\text{-}ss}}$. This concludes the proof.

Proof of Theorem 1 The theorem follows directly from Lemmas 1 and 2.

B DPOR for TSO

In this appendix, we establish the correctness of Theorem 2, which states that the DPOR algorithms *Source-DPOR* and *Optimal-DPOR* of [1], when based on the happens-before relation induced by chronological traces, explore at least one execution per equivalence class induced by Shasha-Snir traces. Theorem 2 also states that *Optimal-DPOR* explores exactly one execution per equivalence class. We also provide more detail on how a DPOR algorithm, such as *Source-DPOR* of [1], can be used for SMC on programs running under TSO by computing chronological traces on the fly.

Correctness of *Source-DPOR* also implies that several other DPOR algorithms, e.g., in [13,27] would be correct if based on chronological traces. This is because these algorithms are subsumed by *Source-DPOR* in the sense that the set of executions that are explored by these algorithms in some particular analysis includes the set of executions that could be explored by *Source-DPOR* in some analysis.

Theorem 2. (Correctness of DPOR algorithms) *The algorithms* Source-DPOR *and* Optimal-DPOR *of [1], based on the happens-before relation induced by chronological*

traces, explore at least one execution per equivalence class induced by Shasha-Snir traces. Moreover, Optimal-DPOR *explores exactly one execution per equivalence class.*

Proof The proof of Theorem 2 mainly uses the correctness of *Source-DPOR*, which is proven in [1]. More precisely, in [1] it is proven that *Source-DPOR* is correct whenever it is based on an assignment of happens-before relations to executions, which is *valid*. An assignment of happens-before relations $\rightarrow_\tau$ to executions τ is valid if it satisfies the following natural properties (from [1]).

1. $\rightarrow_\tau$ is a partial order on the events in τ, which is included in $<_\tau$,
2. the events of each thread are totally ordered by $\rightarrow_\tau$,
3. if τ' is a prefix of τ, then $\rightarrow_\tau$ and $\rightarrow_{\tau'}$ are the same on τ'.
4. the assignment of happens-before relations to executions partitions the set of executions into equivalence classes; i.e., if τ' is a linearization of the happens-before relation on τ, then τ' is assigned the same happens-before relation as τ; we use $\simeq$ to denote the corresponding equivalence relation,
5. whenever τ and τ' are equivalent then they end up in the same global program state,
6. for any sequences τ, τ' and τ'', such that $\tau \cdot \tau''$ is an execution, we have $\tau \simeq \tau'$ if and only if $\tau \cdot \tau'' \simeq \tau' \cdot \tau''$, and
7. if $\tau \cdot (p, i, j)$ is an execution, whose last event is performed by thread p, and q, r are different threads, such that (p, i, j) would "happen before" a subsequent event by r but not a subsequent event by q, then (p, i, j) would also "happen before" (r, i'', j'') in the execution $\tau \cdot (p, i, j) \cdot (q, i', j') \cdot (r, i'', j'')$.

A consequence of these definitions is that that if e and e' are two consecutive events in τ with $e \not\rightarrow_\tau e'$, then e and e' can be swapped without affecting the (global) state after the two events.

The theorem can now be proven by establishing that the happens-before assignment induced by chronological traces is valid. Conditions 1, 2, 3, and 6 follow straightforwardly from definitions Condition 4 follows by observing that changing the order between non-related events does not affect the definition of the chronological trace. Condition 5 follows by observing that the chronological trace captures all dependences that are needed for determining which values are read and written by loads and stores. Finally, Condition 7 follows by noting that an arrow between (p, i, j) and (r, i'', j'') in a chronological trace cannot be removed by inserting an event that is independent with p. This concludes the proof of Theorem 2. □

We next provide more details on the computation of the happens-before relation induced by chronological traces.

The happens-before relation $\rightarrow_\tau^{\mathsf{ct}}$ is computed using vector clocks, while taking the particular structure of chronological traces into account. The main difference from computing happens-before relations for sequentially consistent executions (see, e.g., [27]) is that load events which get their value by store forwarding are not immediately synchronized with the vector clock of the memory location. Instead the load is associated with the store buffer entry from which it got its value. The load is then synchronized with the memory location at the time when the store buffer entry is updated to memory.

Formally, we introduce auxiliary configurations, and define the semantics of instructions over them. When exploring an execution, all instructions will be applied

simultaneously to the TSO system configuration (as described in Section A.2) and the auxiliary configuration. Below we need vector clocks. A *vector clock* is a function $C : (\mathsf{TID} \cup \mathsf{AuxTID}) \mapsto \mathbb{N}$. The intuition is that C captures a set of observed events. For every thread p, the first $C(p)$ events by p have been observed. We let $\mathsf{VecClocks} = ((\mathsf{TID} \cup \mathsf{AuxTID}) \mapsto \mathbb{N})$ denote the set of vector clocks.

An *auxiliary configuration* is a triple $(\mathcal{C}, \mathcal{B}, \mathcal{M})$, where

$\mathcal{C} : (\mathsf{TID} \cup \mathsf{AuxTID} \cup \mathsf{Event} \cup \{\bot\}) \mapsto \mathsf{VecClocks}$
maps each (real or auxiliary) thread identifier p to a vector clock representing which parts of the execution have been seen by p. Also, $\mathcal{C}$ maps each event e to the value of $\mathcal{C}(\mathsf{tid}(e))$ at the time immediately after executing e. We fix that $\mathcal{C}(\bot) = (\lambda x.0)$ is a zeroed clock.

$\mathcal{B} : \mathsf{TID} \mapsto (\mathsf{MemLoc} \times \mathsf{Event} \times (\mathsf{Event} \cup \{\bot\}))^*$
maps each real (not auxiliary) thread ID p to a word of letters $(\mathbf{x}, e_s, e_l)$, each of which keeps auxiliary state for the corresponding letter in the store buffer $\mathbb{B}(p)$. Here $\mathbf{x}$ is the accessed memory location, e_s is the store event that produced that letter, and e_l is the latest buffer forwarded load event for which the letter has been the source (if there is no such event then $e_l = \bot$).

$\mathcal{M} : \mathsf{MemLoc} \mapsto ((\mathsf{Event} \cup \{\bot\}) \times 2^{\mathsf{Event}})$
maps each memory location $\mathbf{x}$ to a pair (e_u, E_l), where e_u is the latest update event that accessed $\mathbf{x}$ (or $\bot$ if $\mathbf{x}$ has never been updated), and where E_l is a set which for each thread p that has read $\mathbf{x}$ contains the latest event of p that read the value of $\mathbf{x}$.

Initially all clocks in $\mathcal{C}$ are zeroed, all buffers in $\mathcal{B}$ are empty, and for all memory locations $\mathbf{x}$ we have $\mathcal{M}(\mathbf{x}) = (\bot, \varnothing)$.

The idea here is that as we execute memory accesses, we update the vector clock of the executing thread to reflect which new events have been observed.

For example, when we execute an update $e_{\mathbf{x}}$ which corresponds to a buffer entry $(\mathbf{x}, e_s, e_l)$, we look to the memory $\mathcal{M}(\mathbf{x}) = (e_u, E_l)$. We know that the update event is ordered after the previous update e_u, as well as the previous loads in E_l and the store event e_s which enabled the update e_x. We update the vector clock $\mathcal{C}(\mathsf{tid}(e_{\mathbf{x}}))$ of the auxiliary thread to include all these newly observed events.

The procedure for a load from memory is similar, except that we do not observe previous loads. More interesting are loads that are satisfied by buffer forwarding. When we execute a buffer forwarded load e_l to $\mathbf{x}$, we do not observe *any* new event, since the load was not able to reach and synchronize with the memory. Instead we save the load event with the buffer entry from which it read its value. When that entry is updated to memory, by the update event $\mathsf{upd}_{\mathsf{ld}}(e_l)$, we move e_l to the set of loads that have been observed by $\mathcal{M}(\mathbf{x})$. By this scheme the load event e_l becomes available for observation by precisely the update events which succeed $\mathsf{upd}_{\mathsf{ld}}(e_l)$. In the remainder of this section we will make this intuition formal.

We will also need some notation for dealing with vector clocks. For two vector clocks v, v' we use the notation $v + v'$ to denote the vector clock v'' such that $v''(p) = max(v(p), v'(p))$ for all p. For two vector clocks v, v' we say that $v \leq v'$ when $v(p) \leq v'(p)$ for all p. We say that $v < v'$ if at least one of the inequalities is strict. For an event e and a set E of events we define $E \oplus e = \{e' \in E | \mathsf{tid}(e') \neq \mathsf{tid}(e)\} \cup \{e\}$, i.e. $E \oplus e$ is E where e replaces the previous event $e' \in E$ s.t. $\mathsf{tid}(e') = \mathsf{tid}(e)$. We use the

shorthand $f[x_0, x_1, \cdots, x_n \hookleftarrow v]$ to denote $f[x_0 \hookleftarrow v][x_1 \hookleftarrow v] \cdots [x_n \hookleftarrow v]$, i.e., an assignment of the same value to multiple function arguments.

For two arbitrary auxiliary configurations $c = (\mathcal{C}, \mathcal{B}, \mathcal{M})$ and $c' = (\mathcal{C}', \mathcal{B}', \mathcal{M}')$ we now define the transition relation $c \xrightarrow{p} c'$ depending on the next instruction op of p. We let $j = \mathcal{C}(p)(p) + 1$ be the index of the next event for p and $C_p = \mathcal{C}(p)[p \hookleftarrow j]$ be the corresponding clock and $e = (p, op, j)$ be the event itself.

At the same time as we compute the next auxiliary configuration, we also compute the set $R(e)$ of races (e_r, e'_r) such that $e'_r = e$ and e_r is some earlier event. Recall that the races are the pairs of events from different threads related by $\rightarrow_\tau^{\mathsf{uu}}$, $\rightarrow_\tau^{\mathsf{src\text{-}ct}}$, or $\rightarrow_\tau^{\mathsf{cf\text{-}ct}}$.

Store: If $op = \mathsf{st}(\mathbf{x})$, then $c \xrightarrow{p} c'$ iff $\mathcal{C}' = \mathcal{C}[p, e \hookleftarrow C_p]$, and $\mathcal{M}' = \mathcal{M}$, and $\mathcal{B}' = \mathcal{B}[p \hookleftarrow \mathcal{B}(p) \cdot (\mathbf{x}, e, \bot)]$. There are no races: $R(e) = \varnothing$.

Load from memory: If $op = \mathsf{ld}(\mathbf{x})$ and there is no letter on the form $(\mathbf{x}, _, _)$ in $\mathcal{B}(p)$, then $c \xrightarrow{p} c'$ iff $\mathcal{C}' = \mathcal{C}[p, e \hookleftarrow C'_p]$, and $\mathcal{B}' = \mathcal{B}$ and $\mathcal{M}' = \mathcal{M}[\mathbf{x} \hookleftarrow (e_u, E \oplus e)]$, where $\mathcal{M}(\mathbf{x}) = (e_u, E)$. Here $C'_p = C_p + \mathcal{C}(e_u)$ if $e_u \neq \bot$ and $\mathsf{tid}(e_u) \neq \mathsf{upd}(p)$, and $C'_p = C_p$ otherwise. If $\mathcal{C}(e_u) \not\leq C_p$ and $\mathsf{tid}(e_u) \neq \mathsf{upd}(p)$, then we have $R(e) = \{(e_u, e)\}$. Otherwise $R(e) = \varnothing$.

Intuitively, e is ordered after the last update e_u to $\mathbf{x}$, provided that e_u originated in a different thread.

Load from buffer: If $op = \mathsf{ld}(\mathbf{x})$ and $\mathcal{B}(p) = \mathcal{B}_0 \cdot (\mathbf{x}, e_u, e_l) \cdot \mathcal{B}_1$ for some $\mathcal{B}_0, \mathcal{B}_1, e_u, e_l$ with no letters on the form $(\mathbf{x}, _, _)$ in $\mathcal{B}_1$, then $c \xrightarrow{p} c'$ iff $\mathcal{C}' = \mathcal{C}[p, e \hookleftarrow C_p]$ and $\mathcal{B}' = \mathcal{B}_0 \cdot (\mathbf{x}, e_u, e) \cdot \mathcal{B}_1$ and $\mathcal{M}' = \mathcal{M}$.

Notice that e replaces e_l in the store buffer entry. There are no races: $R(e) = \varnothing$.

Fence: If $op = \mathsf{fence}$ then $c \xrightarrow{p} c'$ iff $\mathcal{B}' = \mathcal{B}$ and $\mathcal{M}' = \mathcal{M}$ and $\mathcal{C}' = \mathcal{C}[p, e \hookleftarrow C_p + \mathcal{C}(\mathsf{upd}(p))]$. There are no races: $R(e) = \varnothing$.

Notice that the semantics of the fence, as defined in Section A.2, guarantee that $\mathcal{B}(p) = \varepsilon$. Hence the vector clock of the auxiliary thread $\mathsf{upd}(p)$ includes the clocks of all updates of $\mathsf{upd}(p)$ corresponding to earlier stores. So e will be ordered after all updates of $\mathsf{upd}(p)$, as prescribed by $\rightarrow_\tau^{\mathsf{uf}}$ for chronological traces.

Update: If $op = \mathsf{u}(\mathbf{x})$ then $c \xrightarrow{p} c'$ iff $\mathcal{B} = (\mathbf{x}, e_s, e_r) \cdot \mathcal{B}'$ and $\mathcal{C}' = \mathcal{C}[p, e \hookleftarrow C_p + \mathcal{C}(e_s) + \mathcal{C}(e_u) + \sum_{e_l \in E \text{ s.t. } \mathsf{upd}(\mathsf{tid}(e_l)) \neq p} \mathcal{C}(e_l)]$ where $\mathcal{M}(\mathbf{x}) = (e_u, E)$ and $\mathcal{M}' = \mathcal{M}[\mathbf{x} \hookleftarrow (e, E')]$. Here $E' = E$ if $e_r = \bot$, and $E' = E \oplus e_r$ otherwise. There is a race with every previous access to $\mathbf{x}$ from a different thread:

$$R(e) = \left\{ (e', e) \middle| \begin{array}{l} e' \in E \cup \{e_u\} \wedge e' \neq \bot \wedge \\ \mathsf{tid}(e') \neq p \wedge \mathsf{upd}(\mathsf{tid}(e')) \neq p \wedge \\ \mathcal{C}(e') \not\leq C_p + \mathcal{C}(e_s) \end{array} \right\}$$

C Adaptation for PSO

In this appendix, we show how our techniques can be adapted to the PSO memory model with minor changes. Before we see how to apply our methods to it, we give an informal description of the PSO memory model.

C.1 PSO Semantics

PSO is a strictly more relaxed model than TSO. As described previously, TSO allows reordering of stores with subsequent loads. PSO allows the same reordering, but also allows the reordering of stores with subsequent stores to different memory locations.

This behavior can be explained by an operational semantics similar to the one described in Section A.2 for TSO, but where each thread has a separate store buffer for each memory location. Each store buffer is FIFO-ordered, so stores to the same memory location by the same thread cannot be reordered. But there is no order maintained between stores in different buffers, so stores by the same thread to different locations may update in reversed order.

In Figure 11(a) we give an example of a program where PSO allows more behaviors than TSO. The execution in Figure 11(b) shows how the stores by p to $\mathbf{x}$ and $\mathbf{y}$ update to memory in reversed order. This allows the thread q to read first $\mathbf{y} = 1$ then $\mathbf{x} = 0$, which would be impossible both under SC and TSO.

In the operational semantics for PSO we introduce one auxiliary thread $\mathsf{upd}(p, \mathbf{x})$ for each pair of a thread p and a memory location $\mathbf{x}$. Each such auxiliary thread is responsible for the updates to $\mathbf{x}$ by p, similarly to how $\mathsf{upd}(p)$ under TSO is responsible for all updates of p.

C.2 Chronological Traces for PSO

The adaptation of chronological traces to PSO is straightforward. The following simple adjustment suffices: Since stores from the same thread p to different memory locations $\mathbf{x}$ and $\mathbf{y}$ are updated by different auxiliary threads $\mathsf{upd}(p, \mathbf{x})$ and $\mathsf{upd}(p, \mathbf{y})$, there is no program order edge between the update events for different memory locations under PSO.

A chronological trace for PSO is illustrated in Figure 11(c). Notice that there is no program order edge from $(\mathsf{upd}(p, \mathbf{x}), \mathsf{u}(\mathbf{x}), 1)$ to $(\mathsf{upd}(p, \mathbf{y}), \mathsf{u}(\mathbf{y}), 1)$. Had there been one, the trace would be cyclic.

D Chronological Traces Versus Xtop Objects

In this appendix, we compare chronological traces and Xtop objects [4].

Our goal was to design a mechanism for representing executions under relaxed memory models that would allow a stateless model checking technique to not explore more that one execution per Shasha-Snir trace. In fact, chronological traces allowed us to achieve this goal, since they are acyclic (in contrast to vanilla Shasha-Snir) but still correspond one-to-one to Shasha-Snir traces. This is not the case for Xtop objects [4],

since two Xtop objects may map to the same Shasha-Snir trace. In fact, chronological traces and Xtop objects are different. This is because an Xtop object includes information about which event pairs are reordered (delayed), while a chronological trace does not. As an illustrative example, consider the SB idiom. A chronological trace is given in Fig. 4, and a corresponding Xtop object is given in [4, Fig. 4(b)]. Notice that the edges $f(b) - de- > f(a)$ and $f(c) - se- > f(d)$ in the Xtop object specify that a and b are reordered, while c and d execute in program order. This choice is not imposed by the chronological trace. Indeed, there is a different Xtop object where c and d are reordered instead of a and b. Both of these Xtop objects correspond to the same chronological trace and to the same Shasha-Snir trace (Fig. 3). Hence a POR technique that explores precisely all Xtop objects would unnecessarily explore more objects than our technique.

p	q
store: $\mathbf{x}$:=1	store: $\mathbf{y}$:=1
store: $\mathbf{z}$:=1	store: $\mathbf{z}$:=0
load: \$$r$:=$\mathbf{z}$	load: \$$r$:=$\mathbf{z}$
load: \$$s$:=$\mathbf{y}$	load: \$$s$:=$\mathbf{x}$

(a) A small program illustrating the idiom of Peterson's mutual exclusion algorithm.

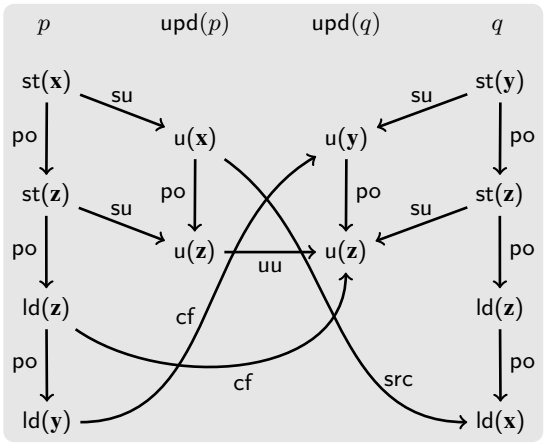

(b) The chronological trace $\mathcal{T}_C(\tau)$ corresponding to the execution in Figure 10(c). Notice that there is no edge between $(q, \mathsf{ld}(\mathbf{z}), 3)$ and either of the updates to $\mathbf{z}$.

(p, st($\mathbf{x}$),1)
  (q, st($\mathbf{y}$),1)
(p, st($\mathbf{z}$),2)
(p, ld($\mathbf{z}$),3)
(p, ld($\mathbf{y}$),4)
  (upd(q), u($\mathbf{y}$),1)
  (q, st($\mathbf{z}$),2)
  (q, ld($\mathbf{z}$),3)
(upd(p), u($\mathbf{x}$),1)
  (q, ld($\mathbf{x}$),4)
(upd(p), u($\mathbf{z}$),2)
  (upd(q), u($\mathbf{z}$),2)

(c) An execution τ.

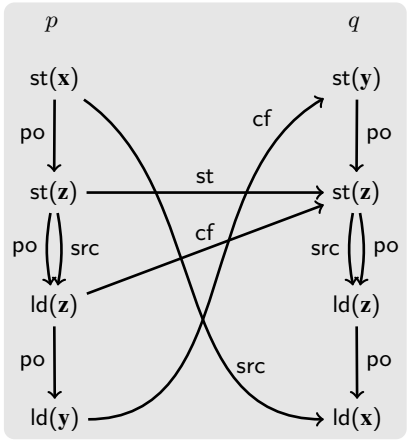

(d) The Shasha-Snir trace $\mathcal{T}(\tau)$ corresponding to the execution in Figure 10(c).

Fig. 10: Traces illustrated by the idiom of Peterson's mutual exclusion algorithm.

p	q
store: $\mathbf{x}$:=1	load: $\$r$:=$\mathbf{y}$
store: $\mathbf{y}$:=1	load: $\$s$:=$\mathbf{x}$

(a) The mp idiom. Possible under PSO: $\$r = 1, \$s = 0$.

$(p, \mathsf{st}(\mathbf{x}), 1)$
$(p, \mathsf{st}(\mathbf{y}), 2)$
$(\mathsf{upd}(p, \mathbf{y}), \mathsf{u}(\mathbf{y}), 1)$
$(q, \mathsf{ld}(\mathbf{y}), 1)$
$(q, \mathsf{ld}(\mathbf{x}), 2)$
$(\mathsf{upd}(p, \mathbf{x}), \mathsf{u}(\mathbf{x}), 1)$

(b) An execution τ where finally $\$r = 1$, $\$s = 0$.

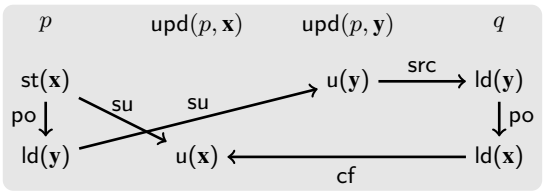

(c) The chronological trace of τ under PSO.

Fig. 11: A behavior allowed under PSO but not under TSO.